LA-UR-19-21268
Approved for public release; distribution is unlimited.

Title: Beyond Moore’s Law: Exploring the Future of Computation

Author(s): Coffrin, Carleton James

Intended for: Invited Seminar

Issued: 2019-02-18

Beyond Moore's Law

Exploring the Future of Computation

Carleton Coffrin

Advanced Network Science Initiative

lanl-ansi.github.io

Who is Carleton?

- **Computer Science / Optimization Algorithms**
  - Los Alamos National Laboratory
  - Ph.D. Advisor Pascal Van Hentenryck

lanl-ansi.github.io

- **Optimization Generalist**
  - Math-heuristics

GRID OPTIMIZATION (GO) COMPETITION

- **Power System Expertise**
  - Convex Power Flow Relaxations (QC formulation)
  - Power Flow Approximations (LPAC formulation)
  - NESTA / PGLib AC-OPF Benchmarking Datasets
  - PowerModels.jl (similar to Matpower, in Julia)

What is Moore’s Law

Shrinking transistors have powered 50 years of advances in computing.

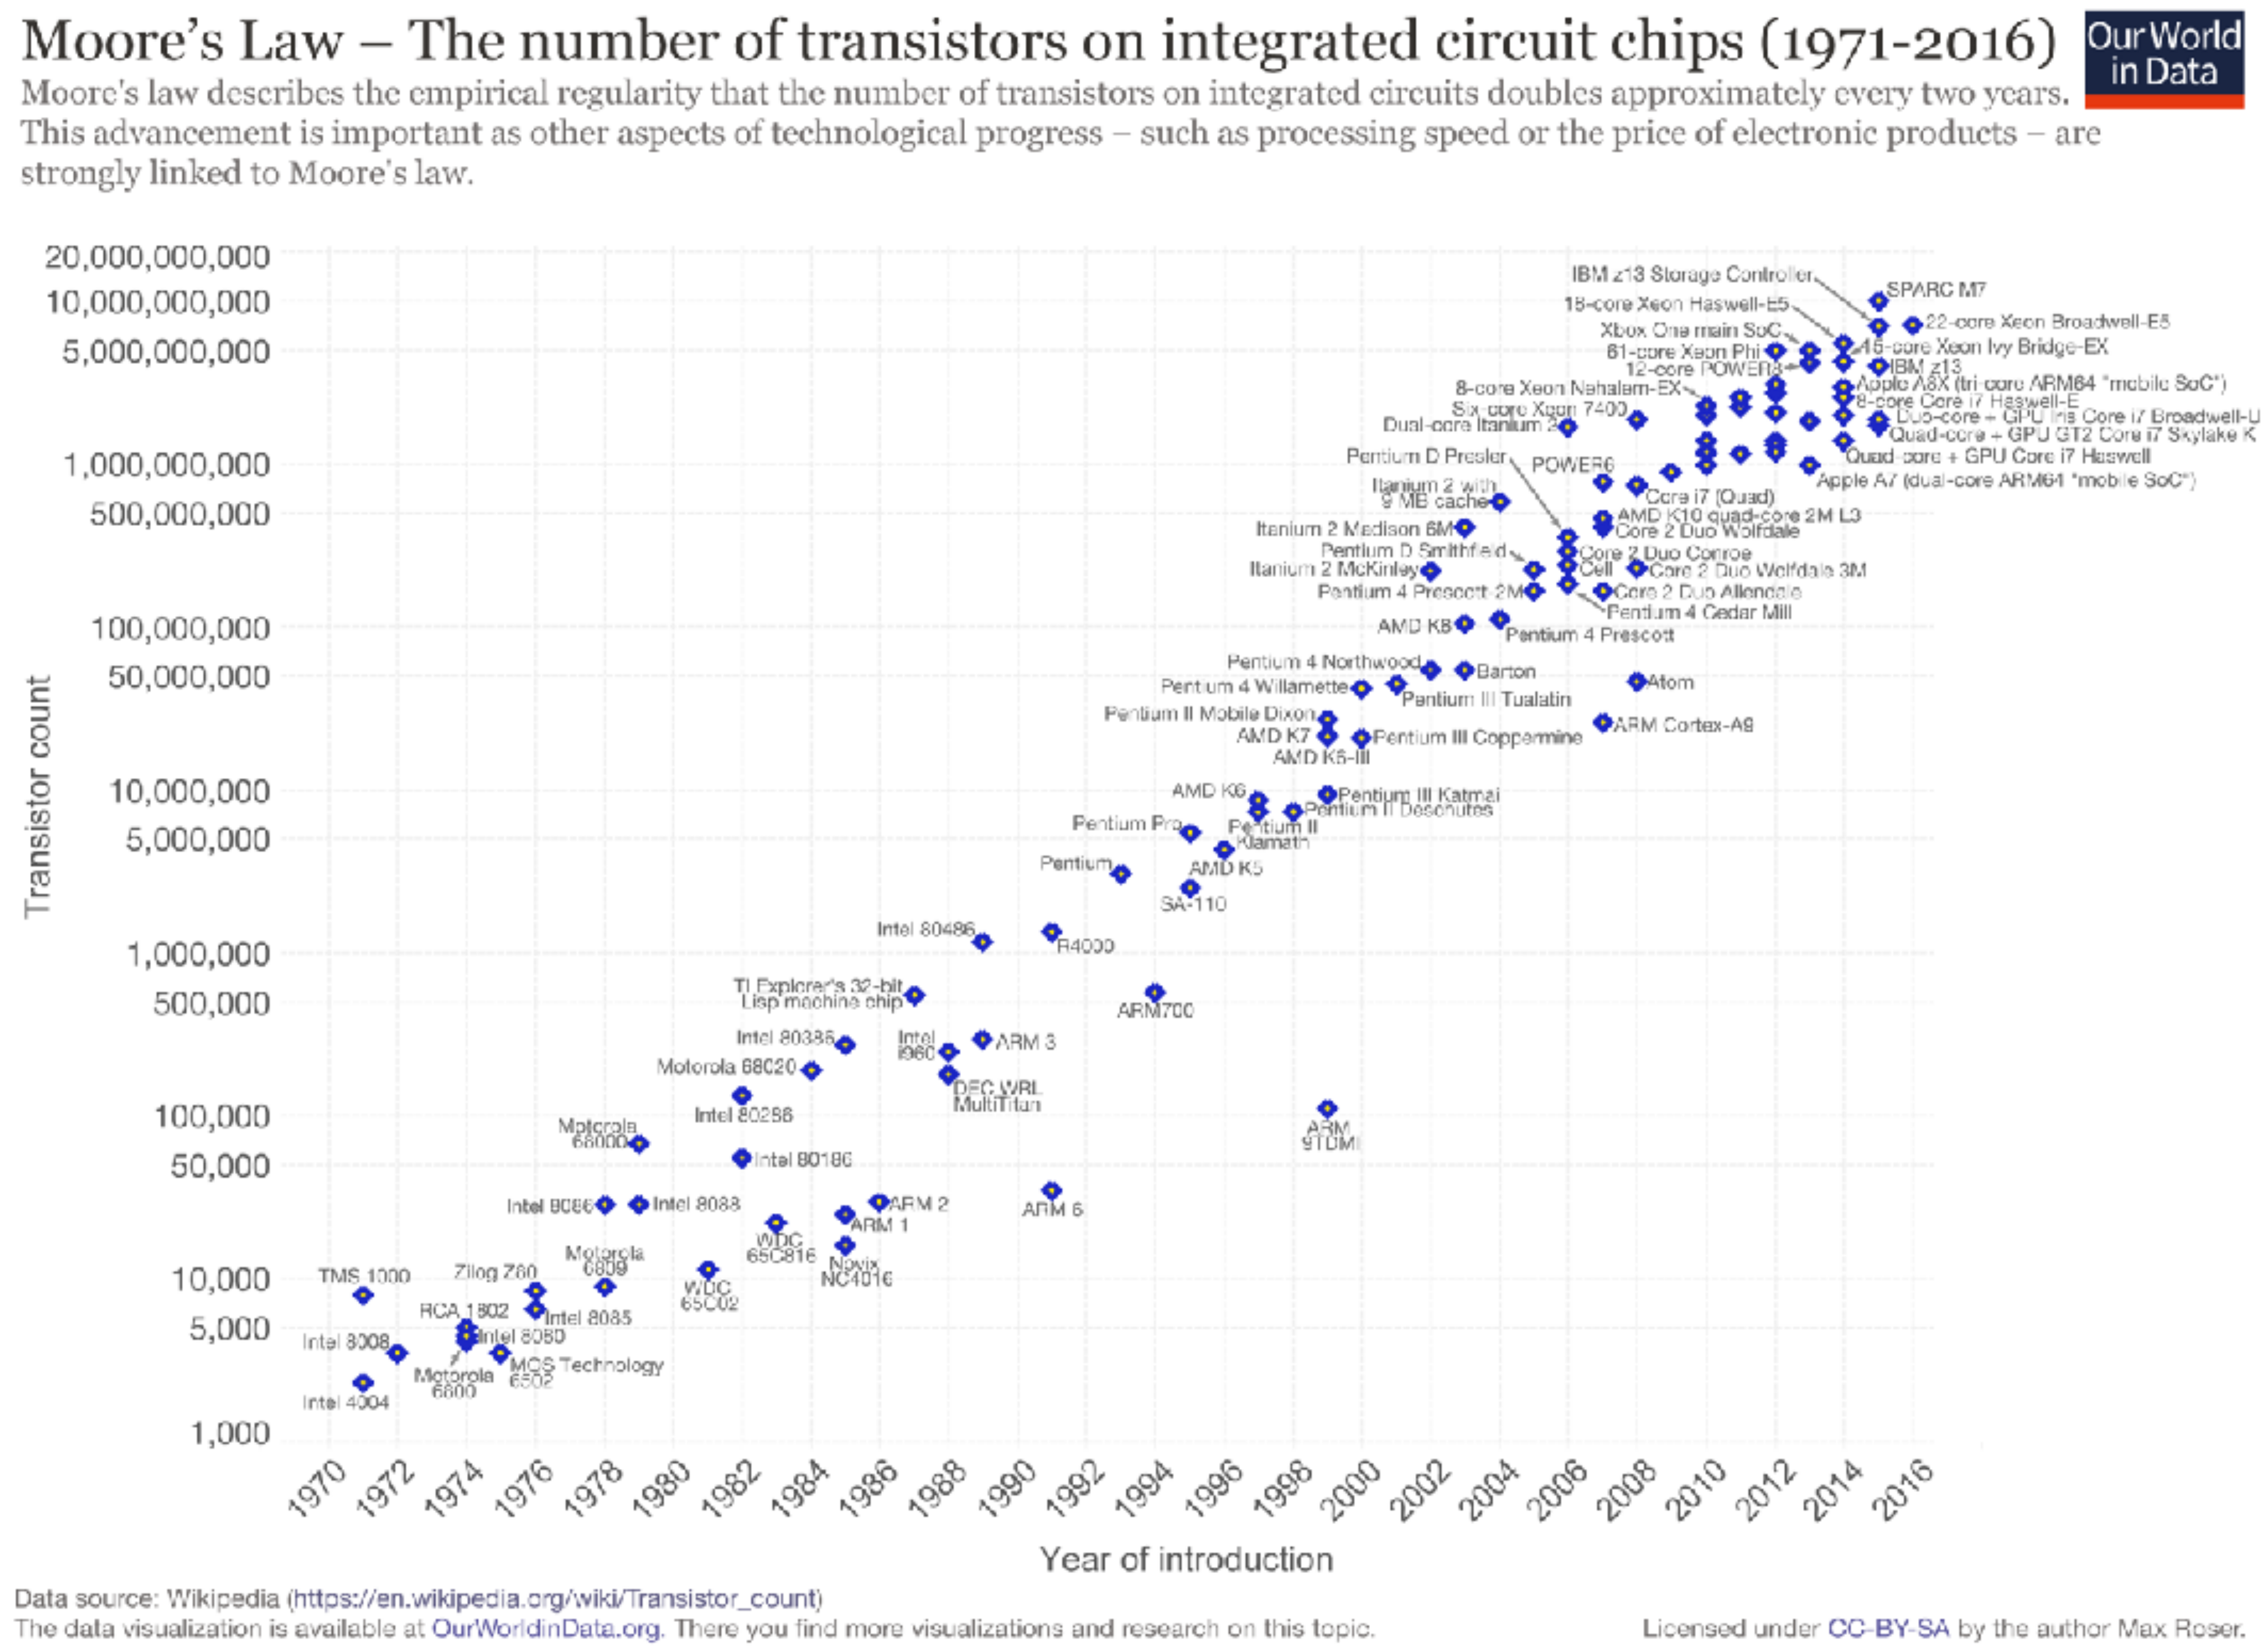

https://commons.wikimedia.org/wiki/File:Moore%27s_Law_Transistor_Count_1971-2016.png

The Bad News

Intel has suggested silicon transistors can only keep shrinking for another five years. (2016)

Intelligent Machines

Intel: Chips Will Have to Sacrifice Speed Gains for Energy Savings

A major technological shift is needed in the next few years if computer chips are to keep improving.

by Katherine Bourzac   February 5, 2016

Hitting basic physical limits, e.g. energy/heat per unit area.

Moore's Law is Dead! Now What?

https://www.technologyreview.com/s/601441/moores-law-is-dead-now-what/

https://www.technologyreview.com/s/600716/intel-chips-will-have-to-sacrifice-speed-gains-for-energy-savings/

Universal / Gate-Based Quantum Computers

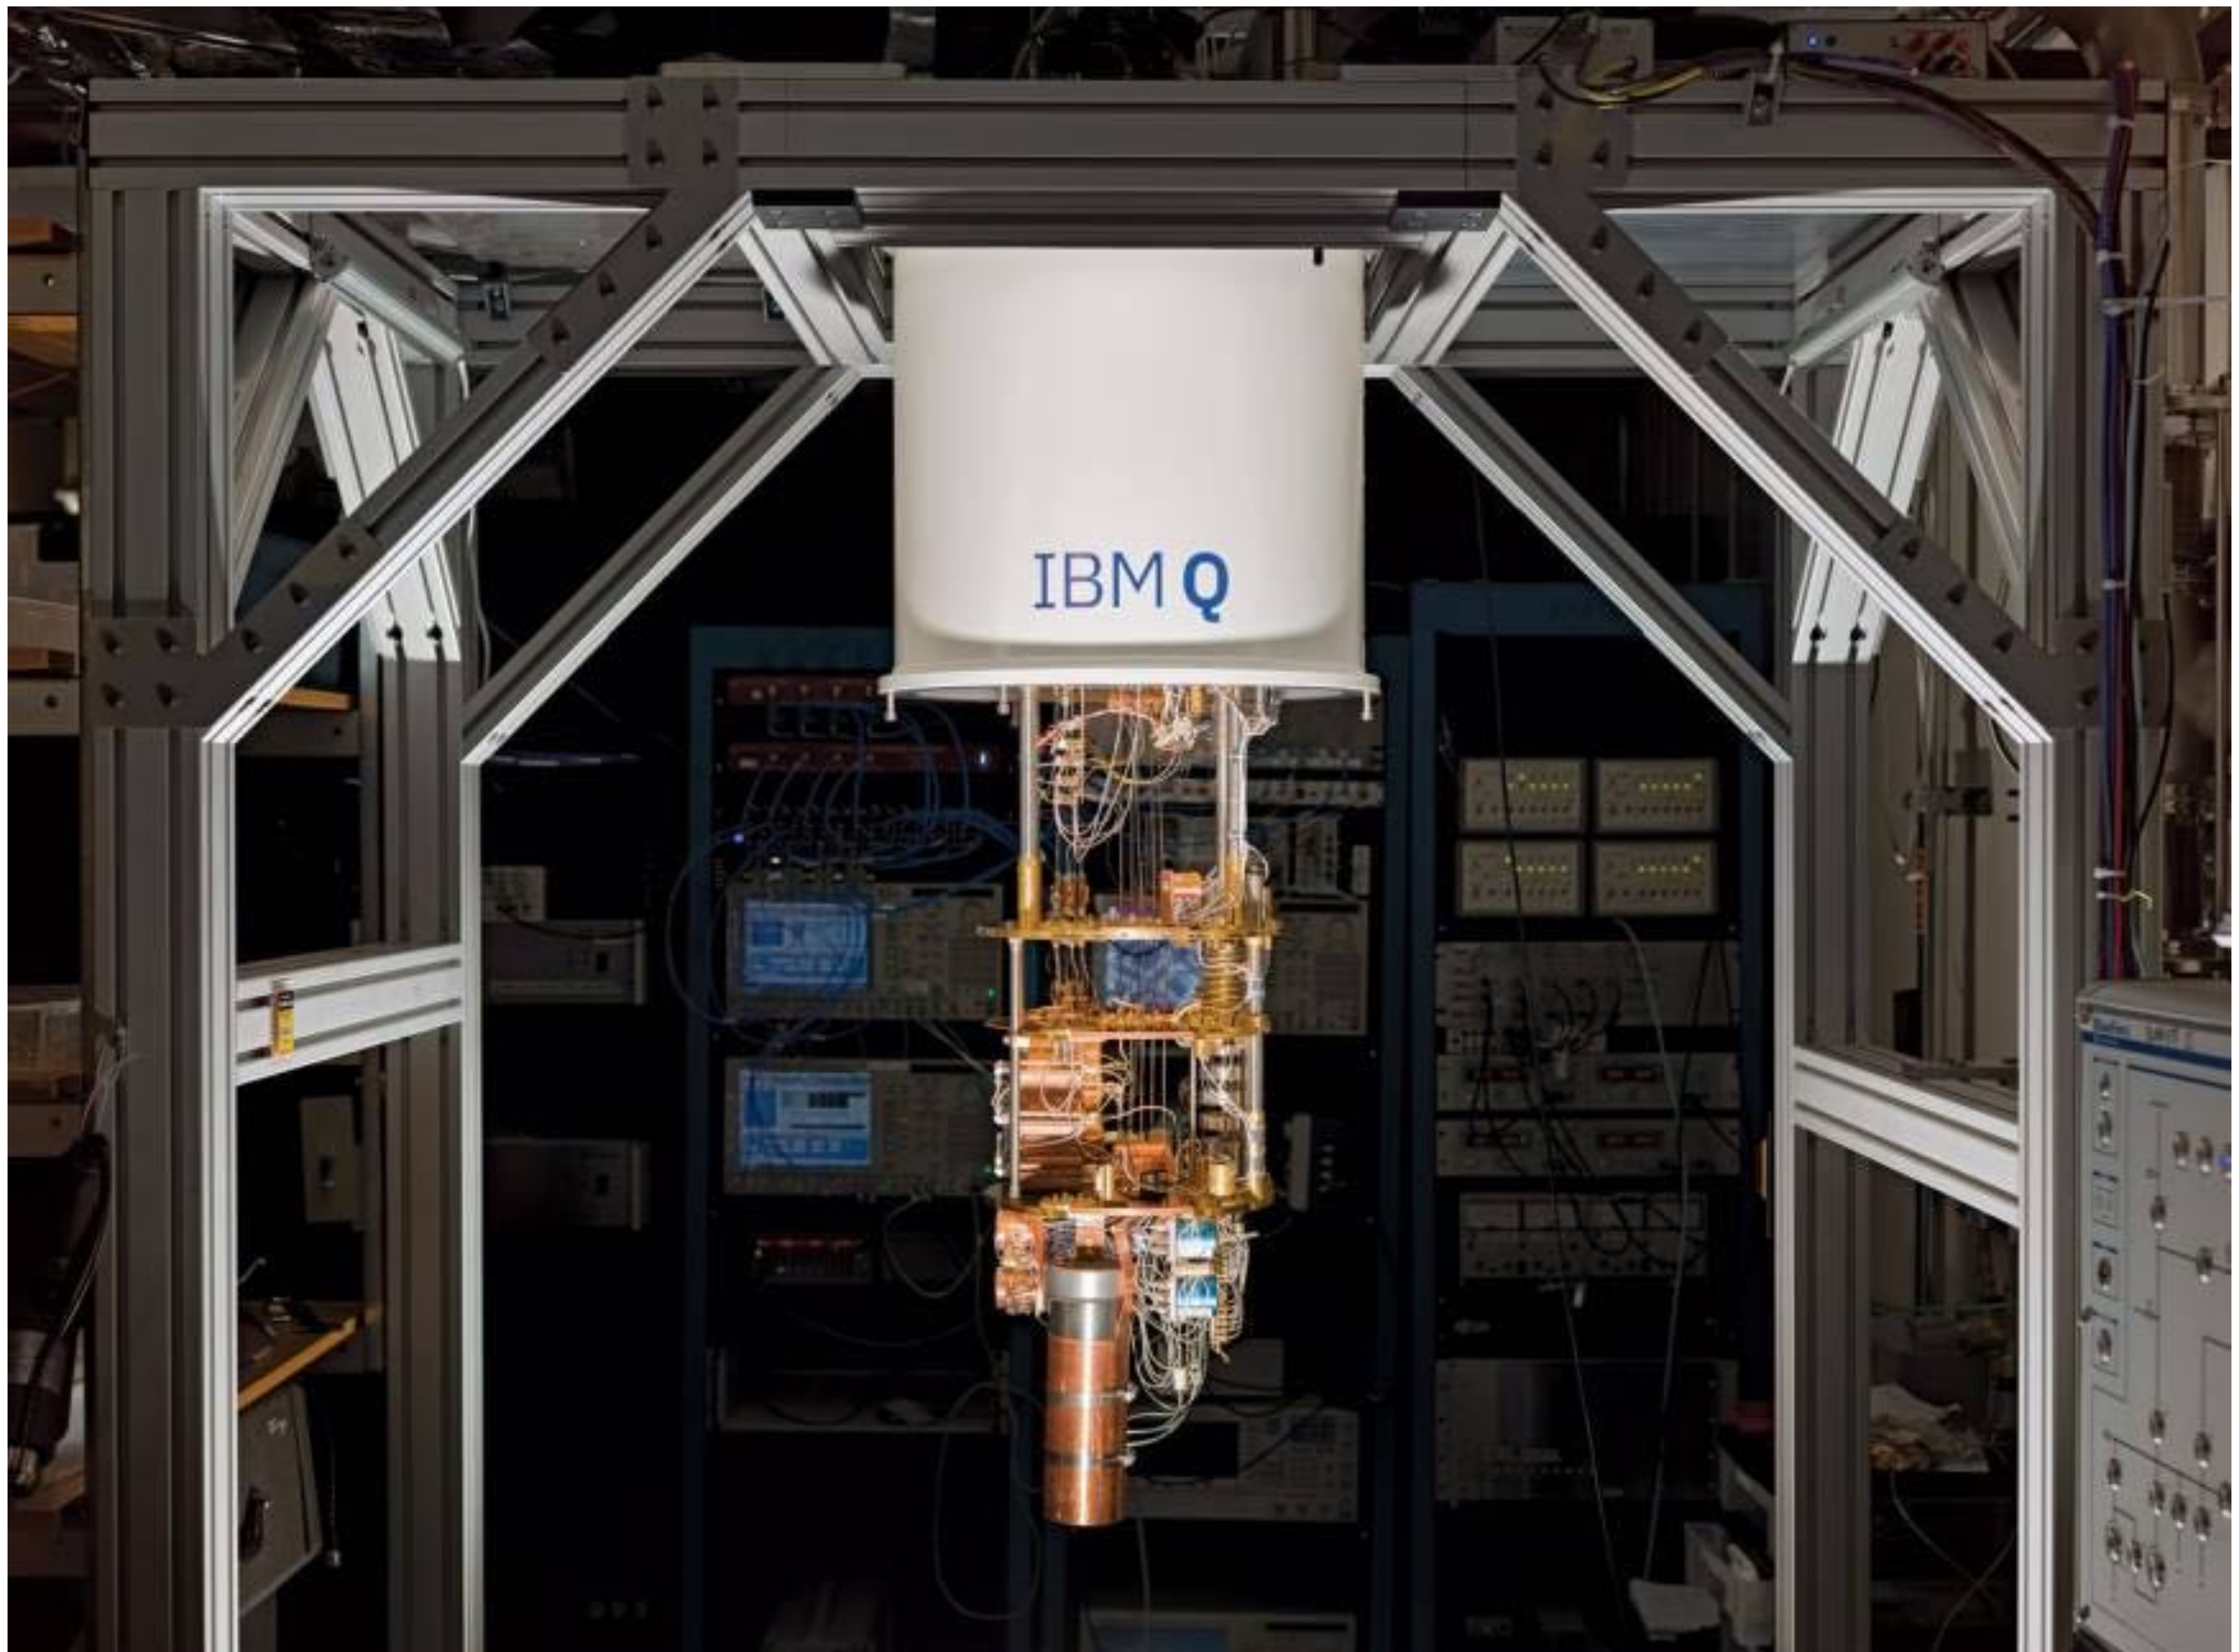

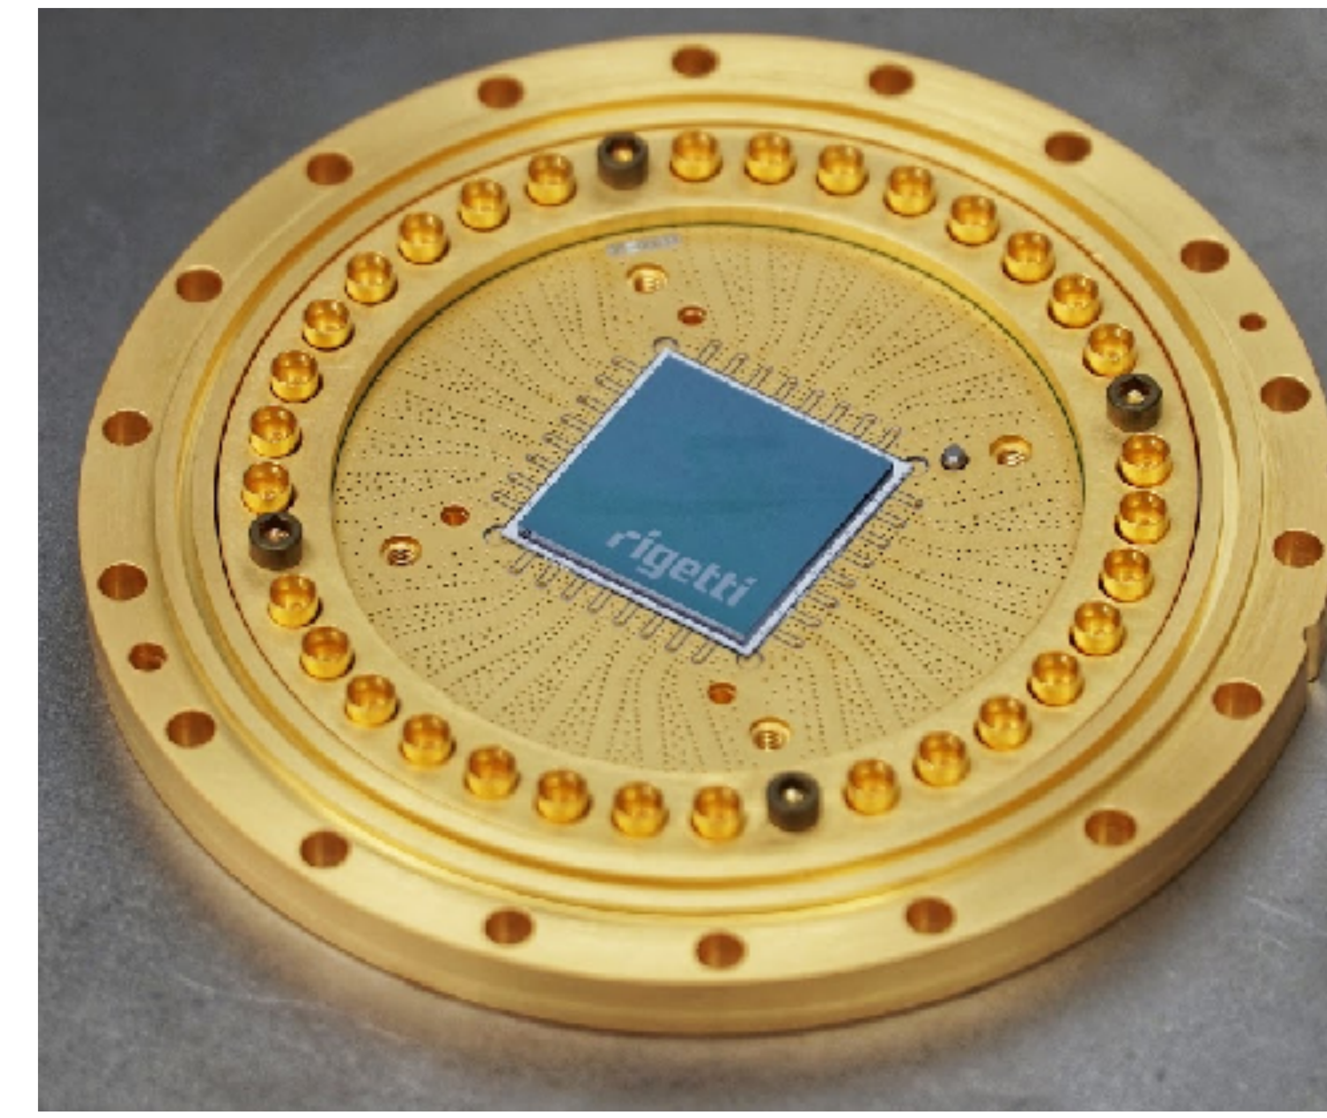

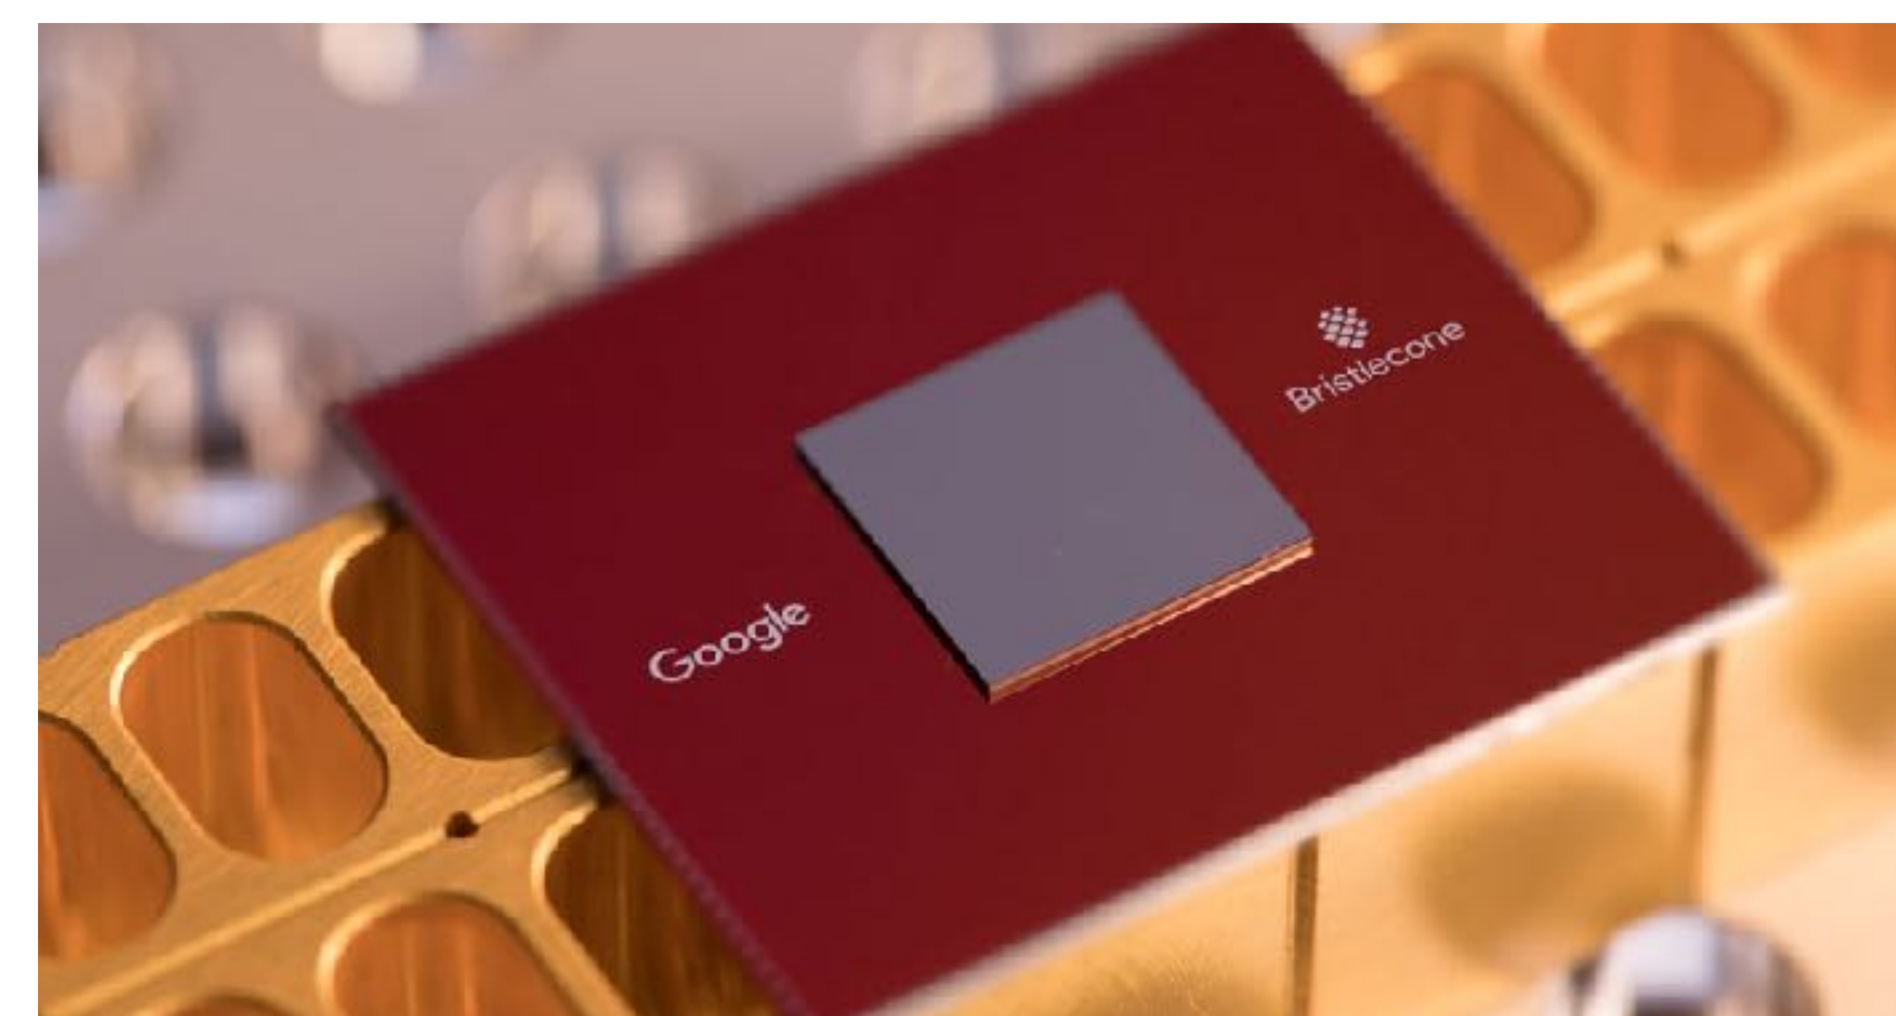

Adiabatic Quantium Computers

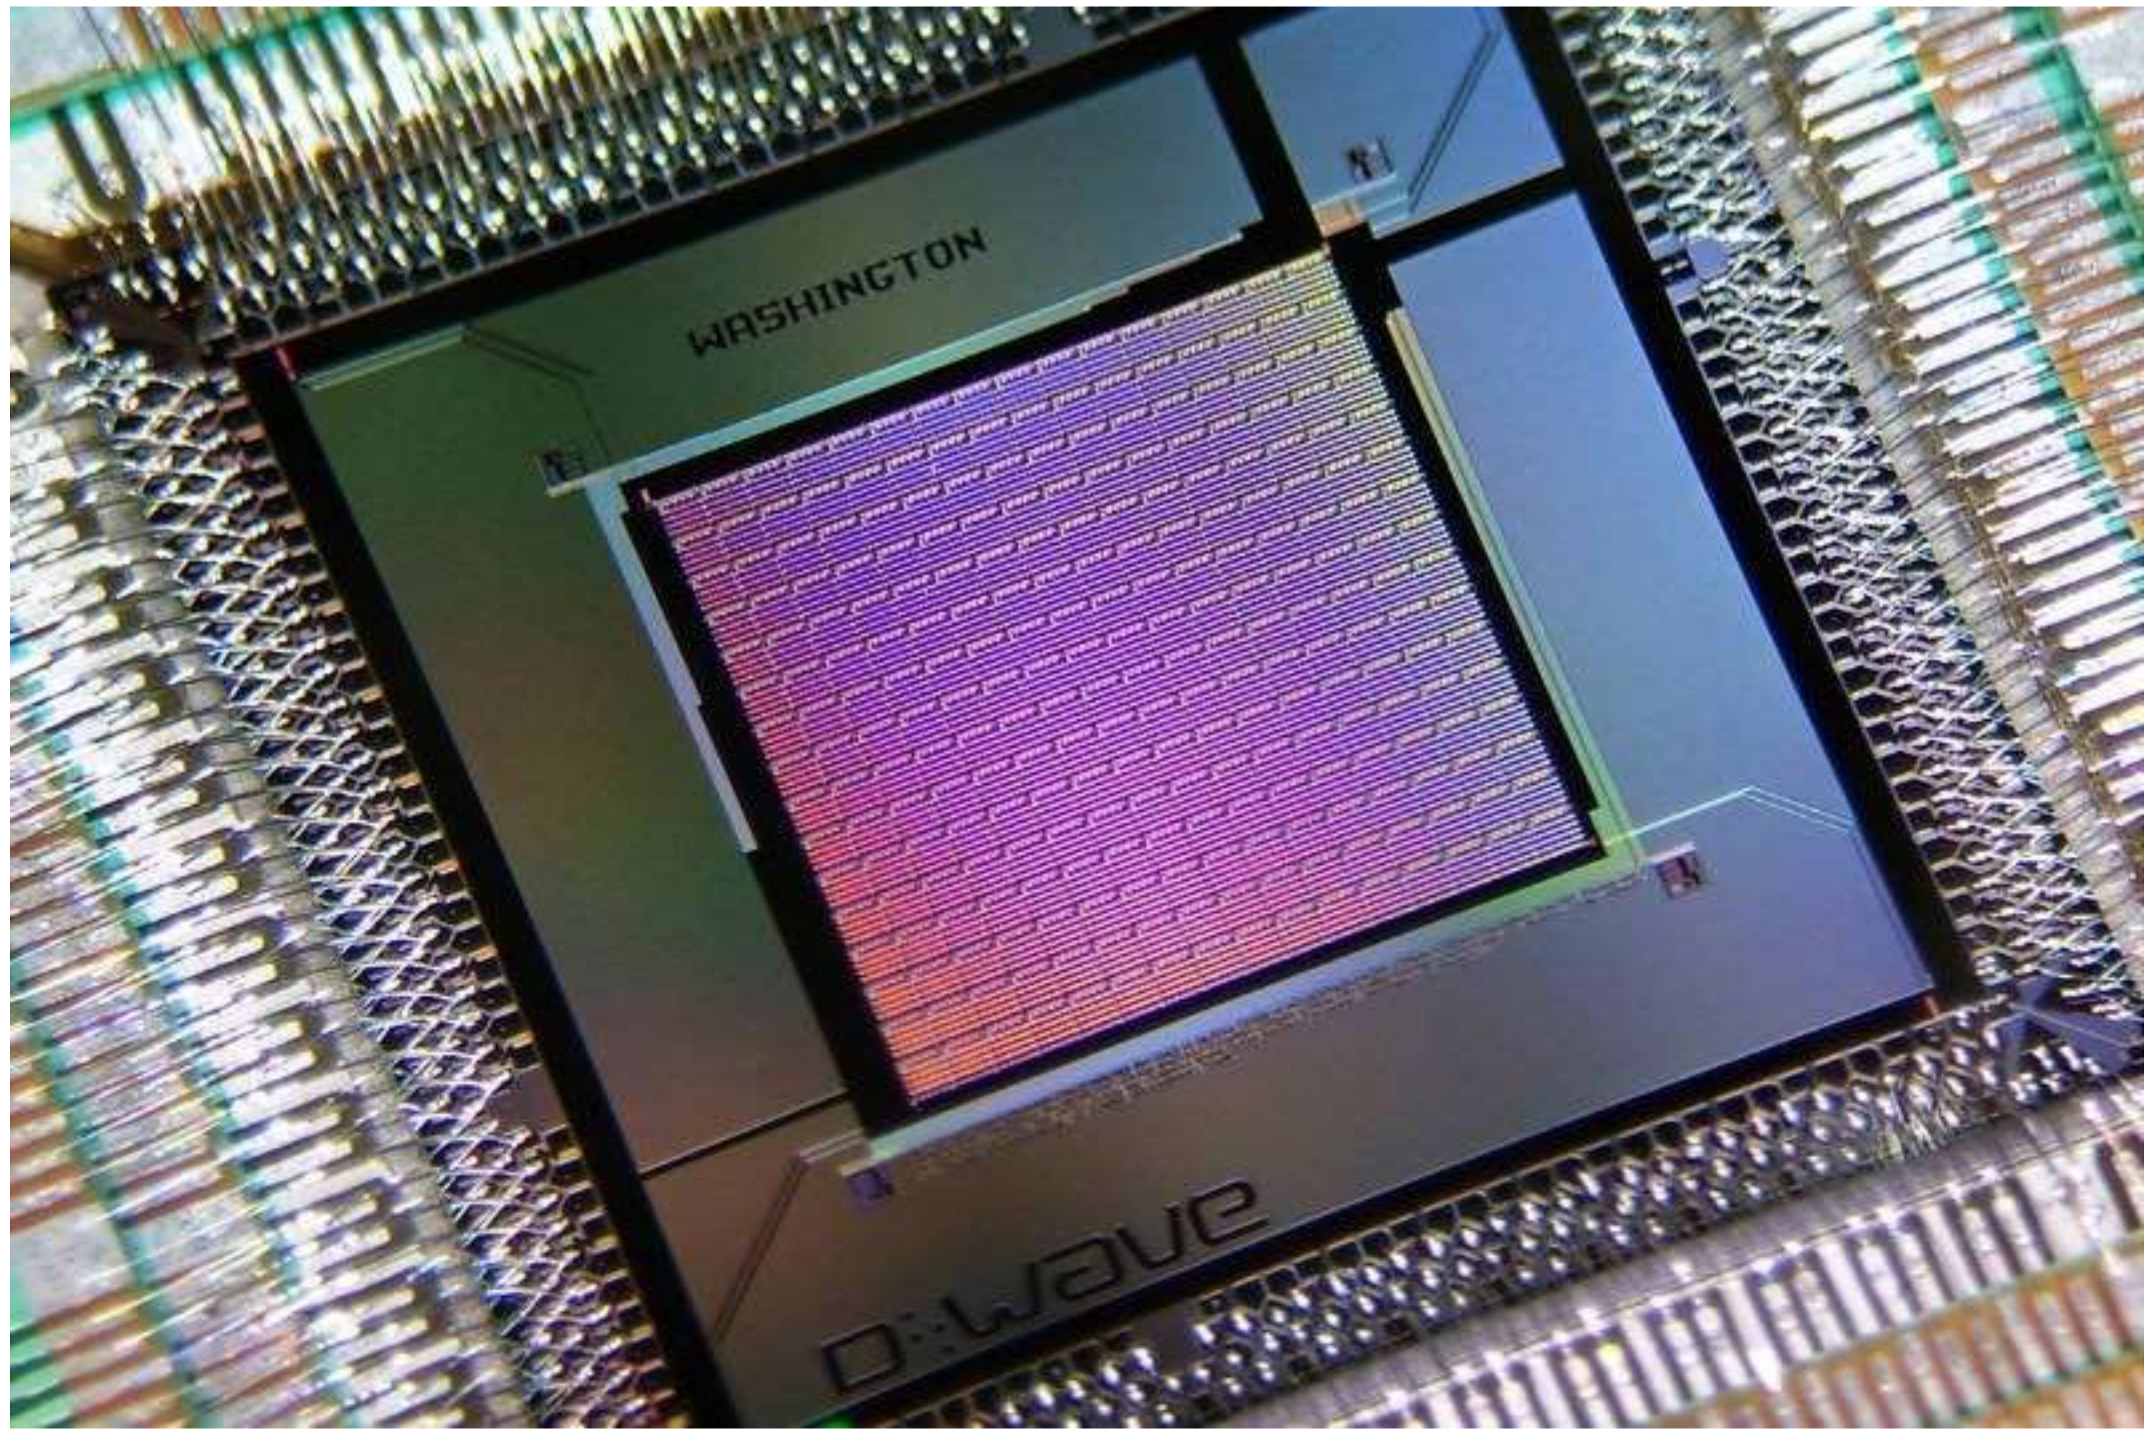

Neuromorphic Coprocessors

IBM True North

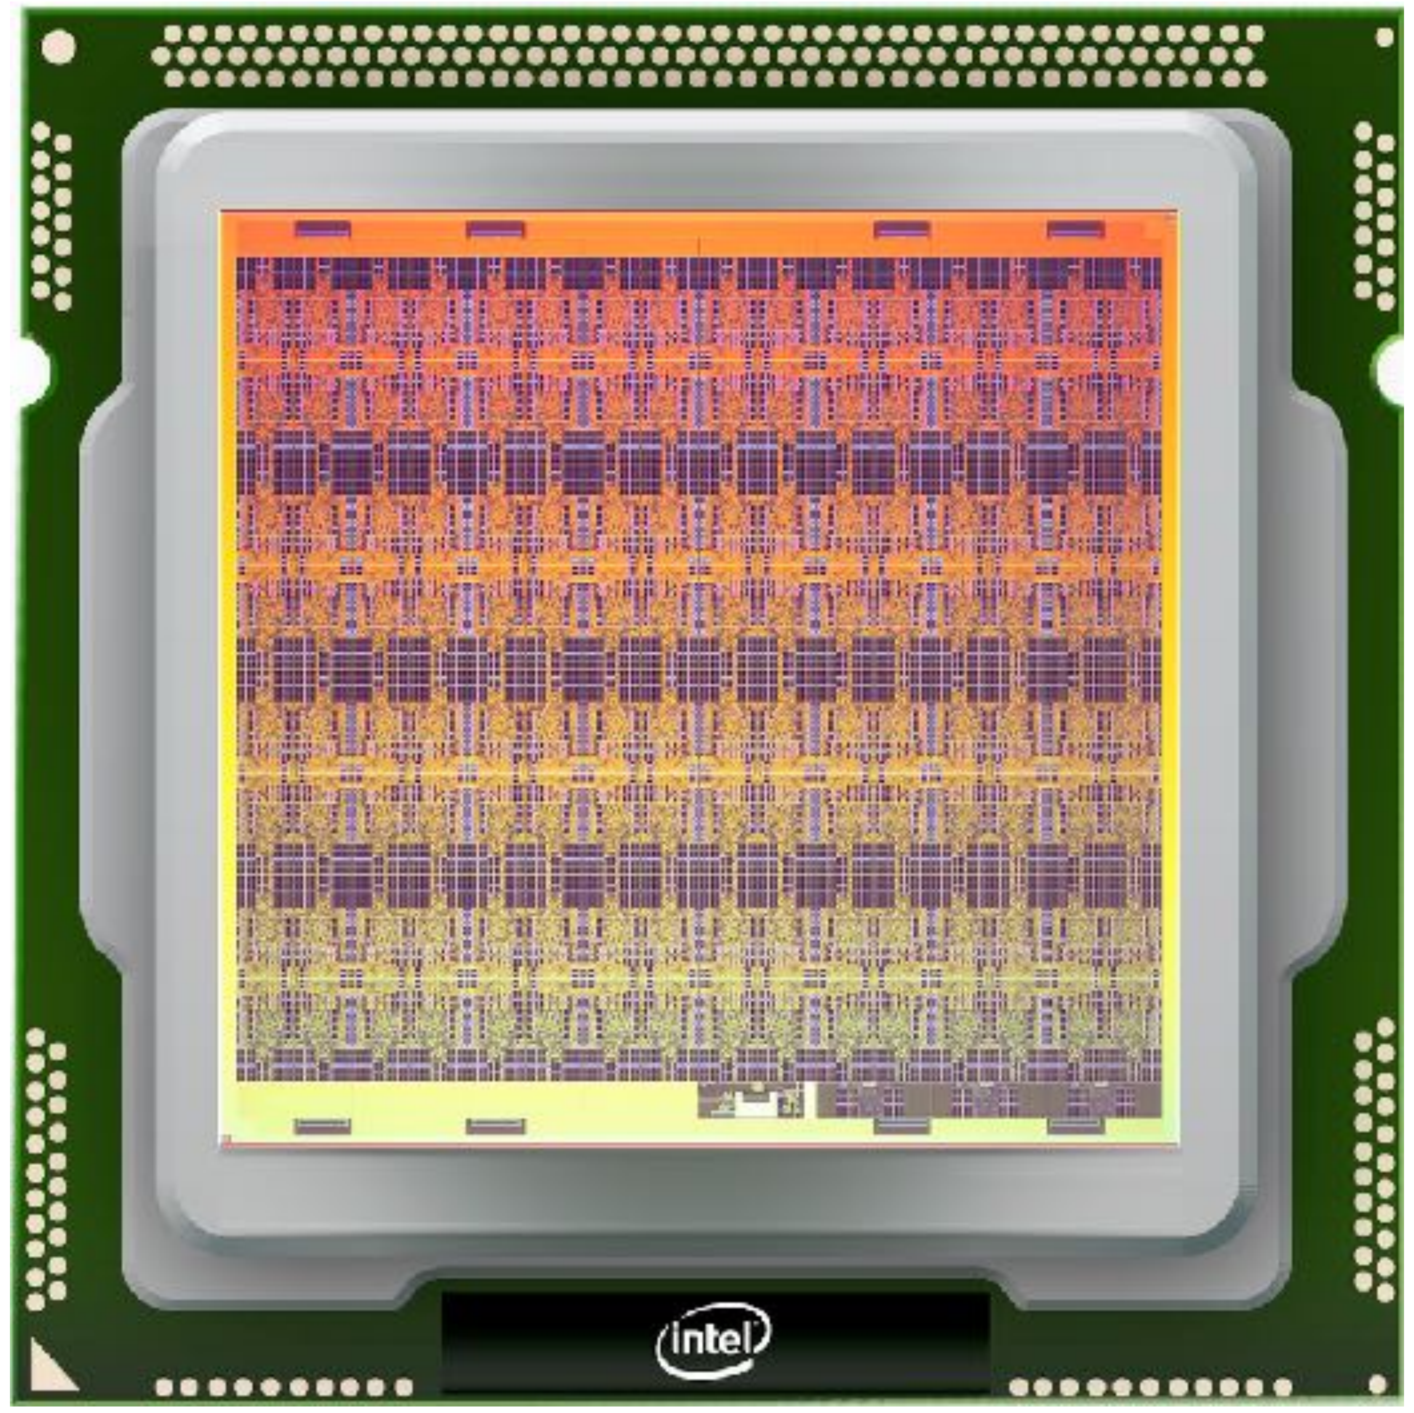

Intel Loihi

Digital Annealers

http://www.fujitsu.com/global/digitalannealer/

Optical Parametric Oscillators

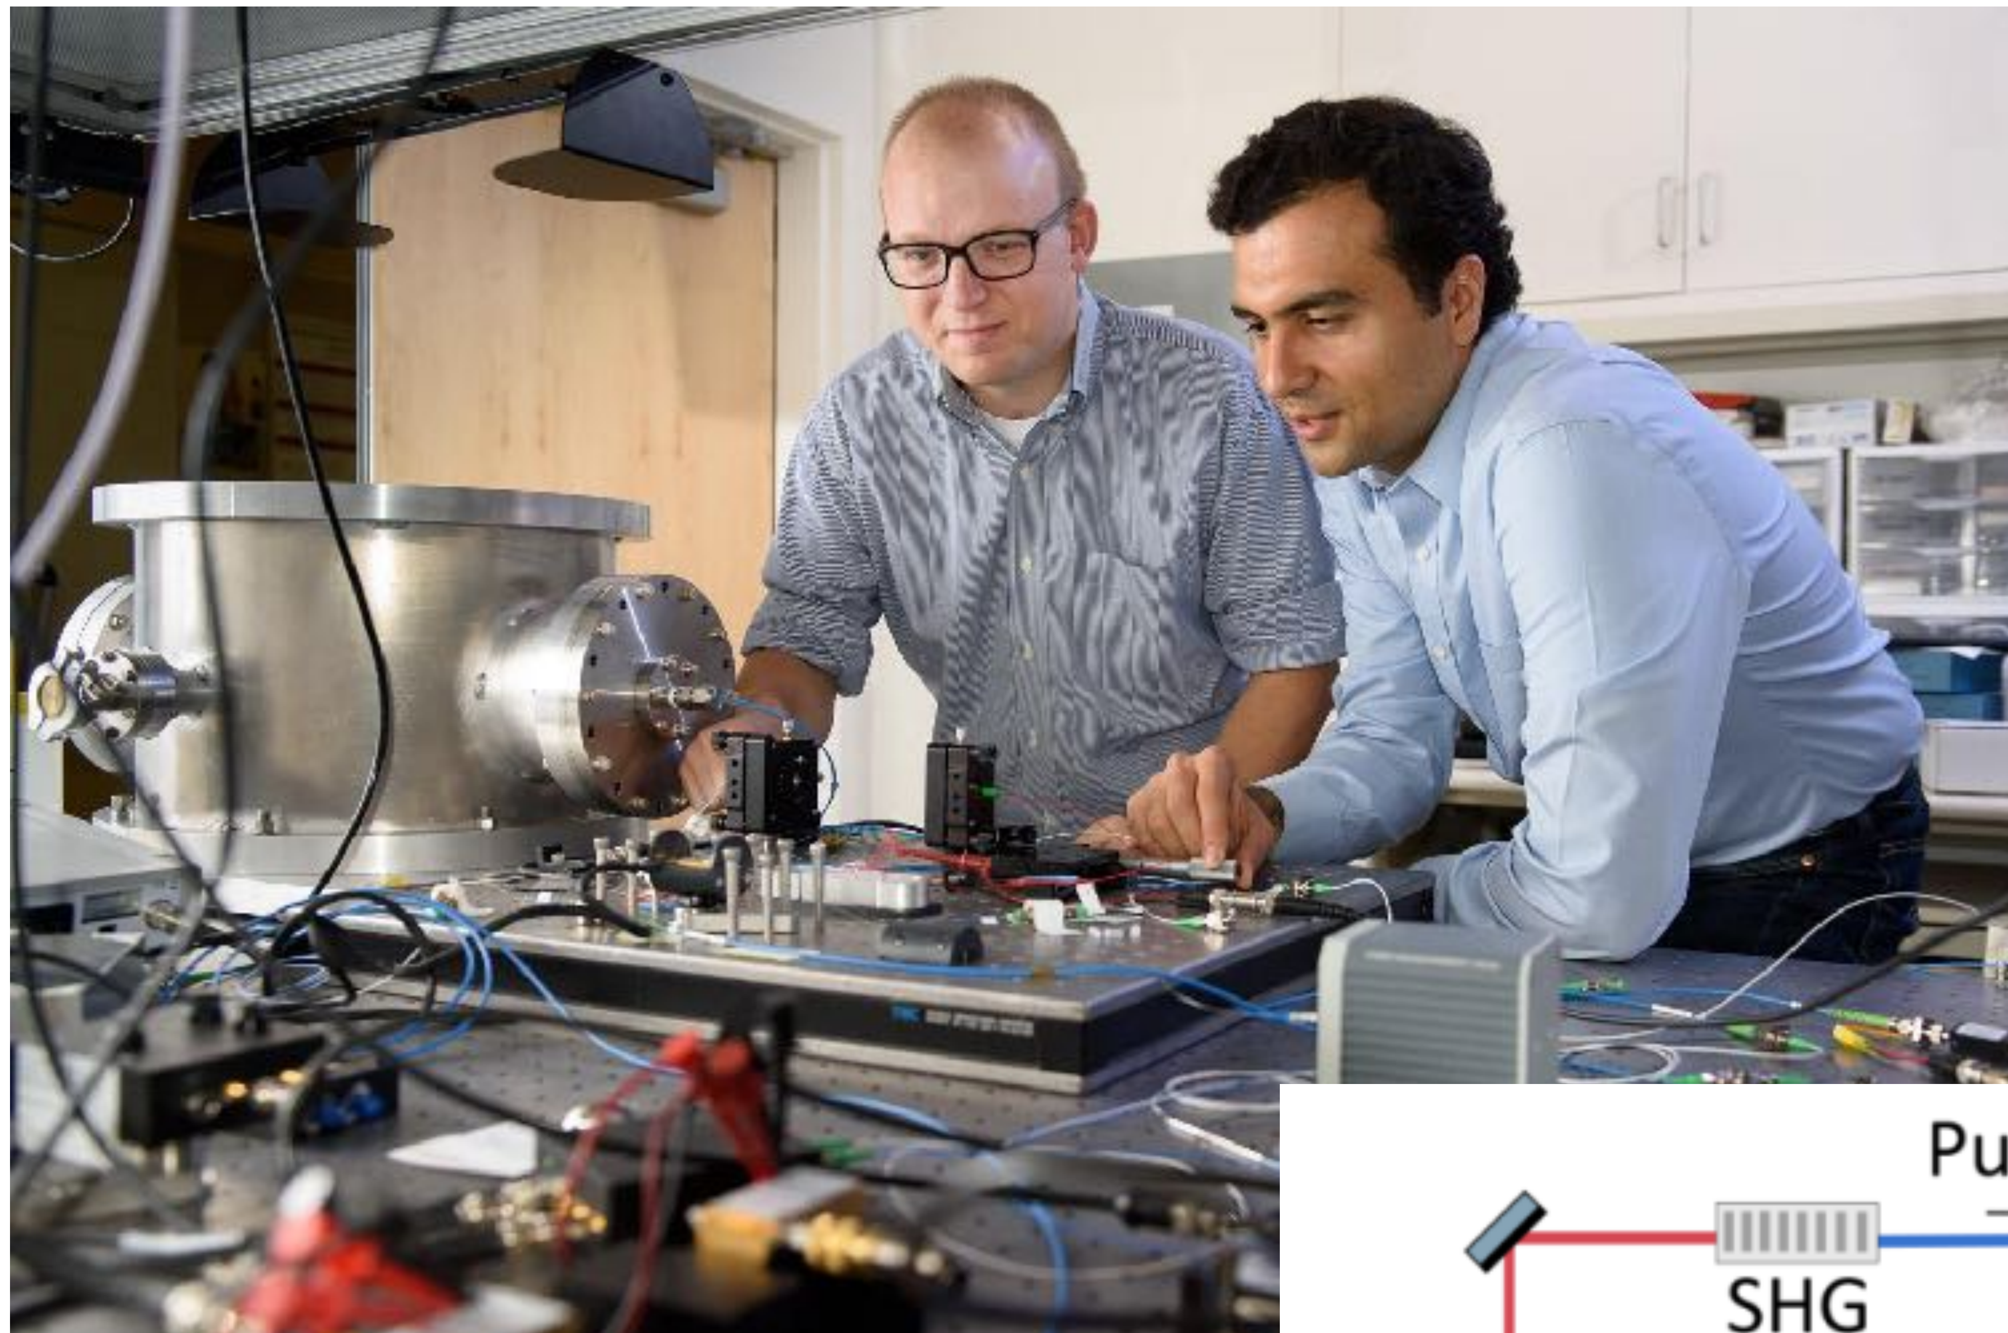

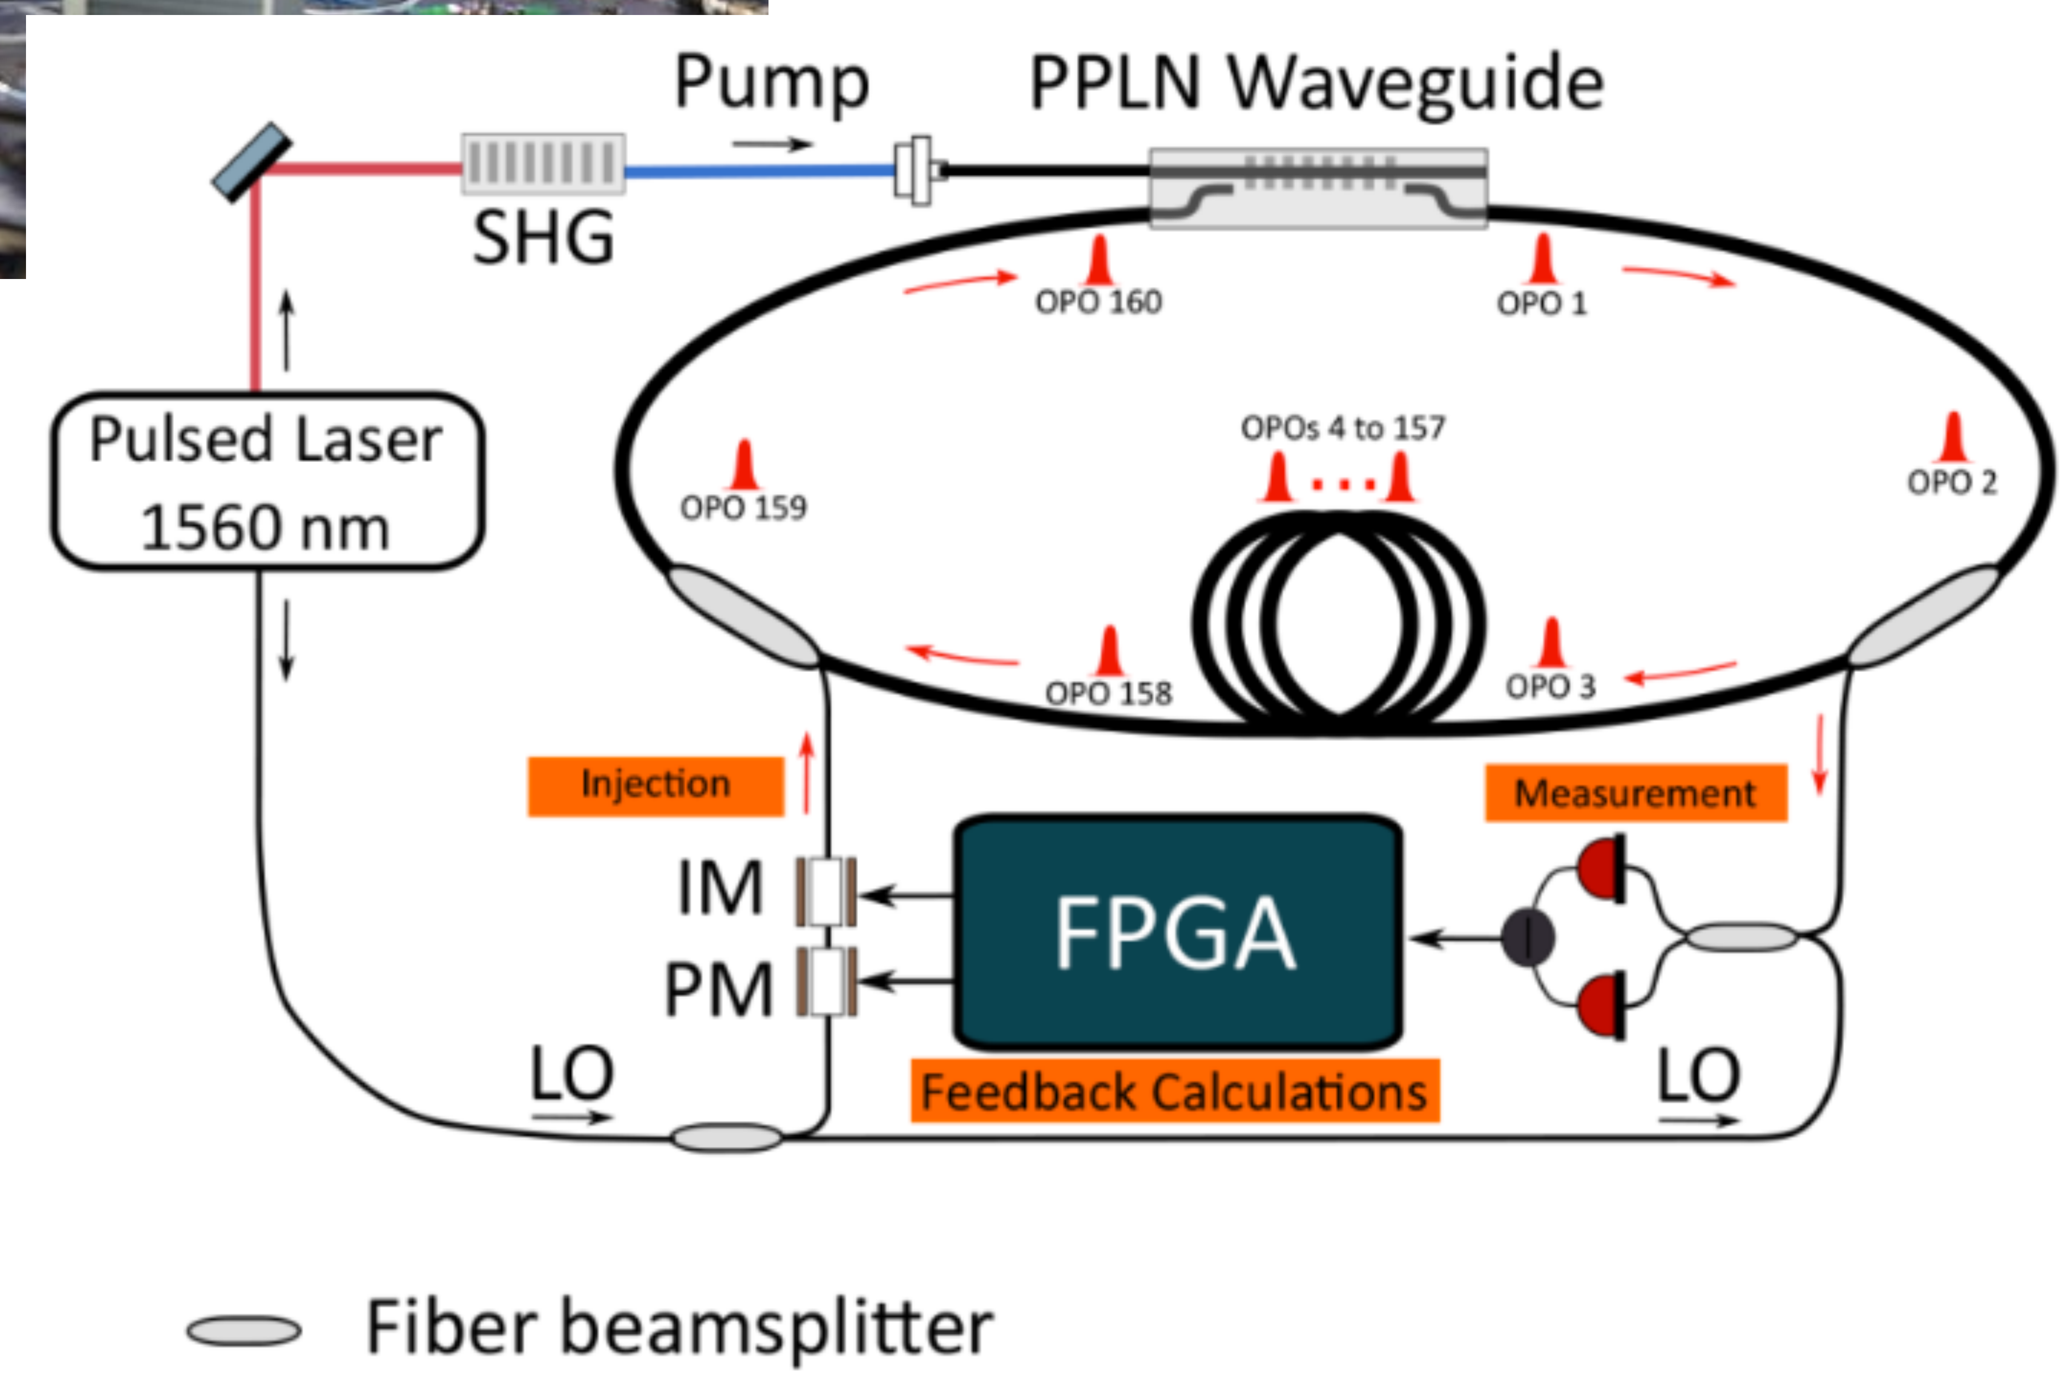

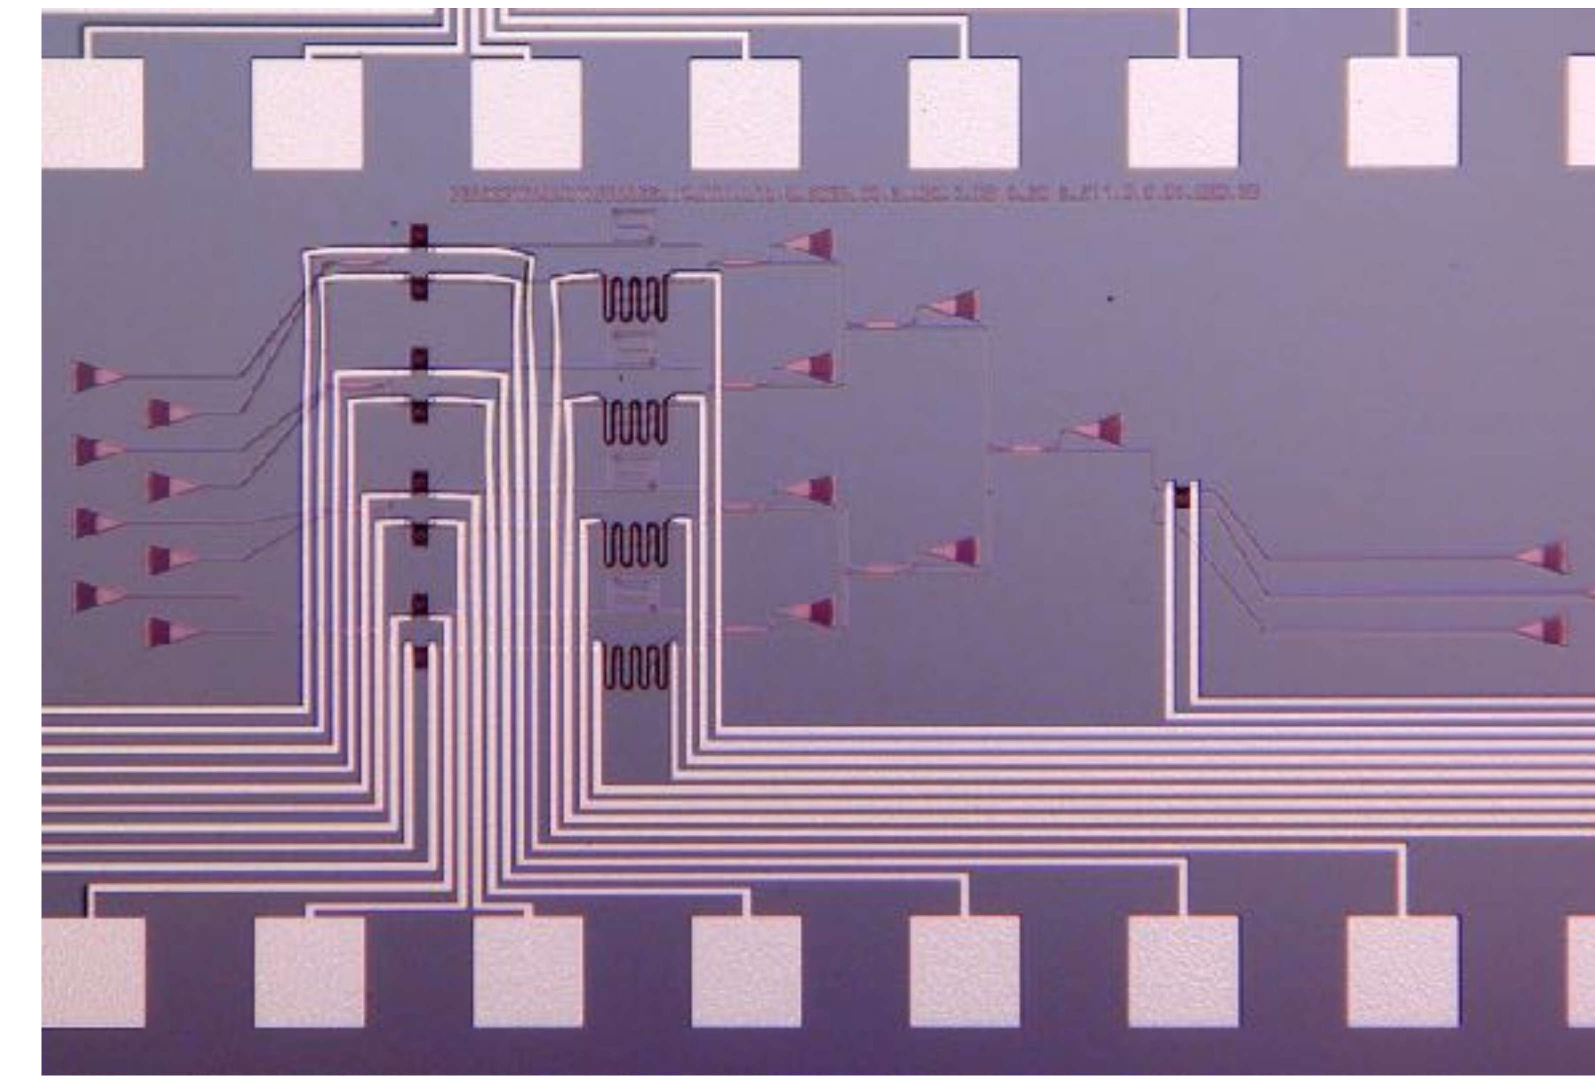

Close-up view of a complex optical circuit created by Hewlett Packard Labs.

Hewlett Packard Labs

https://www.labs.hpe.com/next-next/light

Memristor Networks

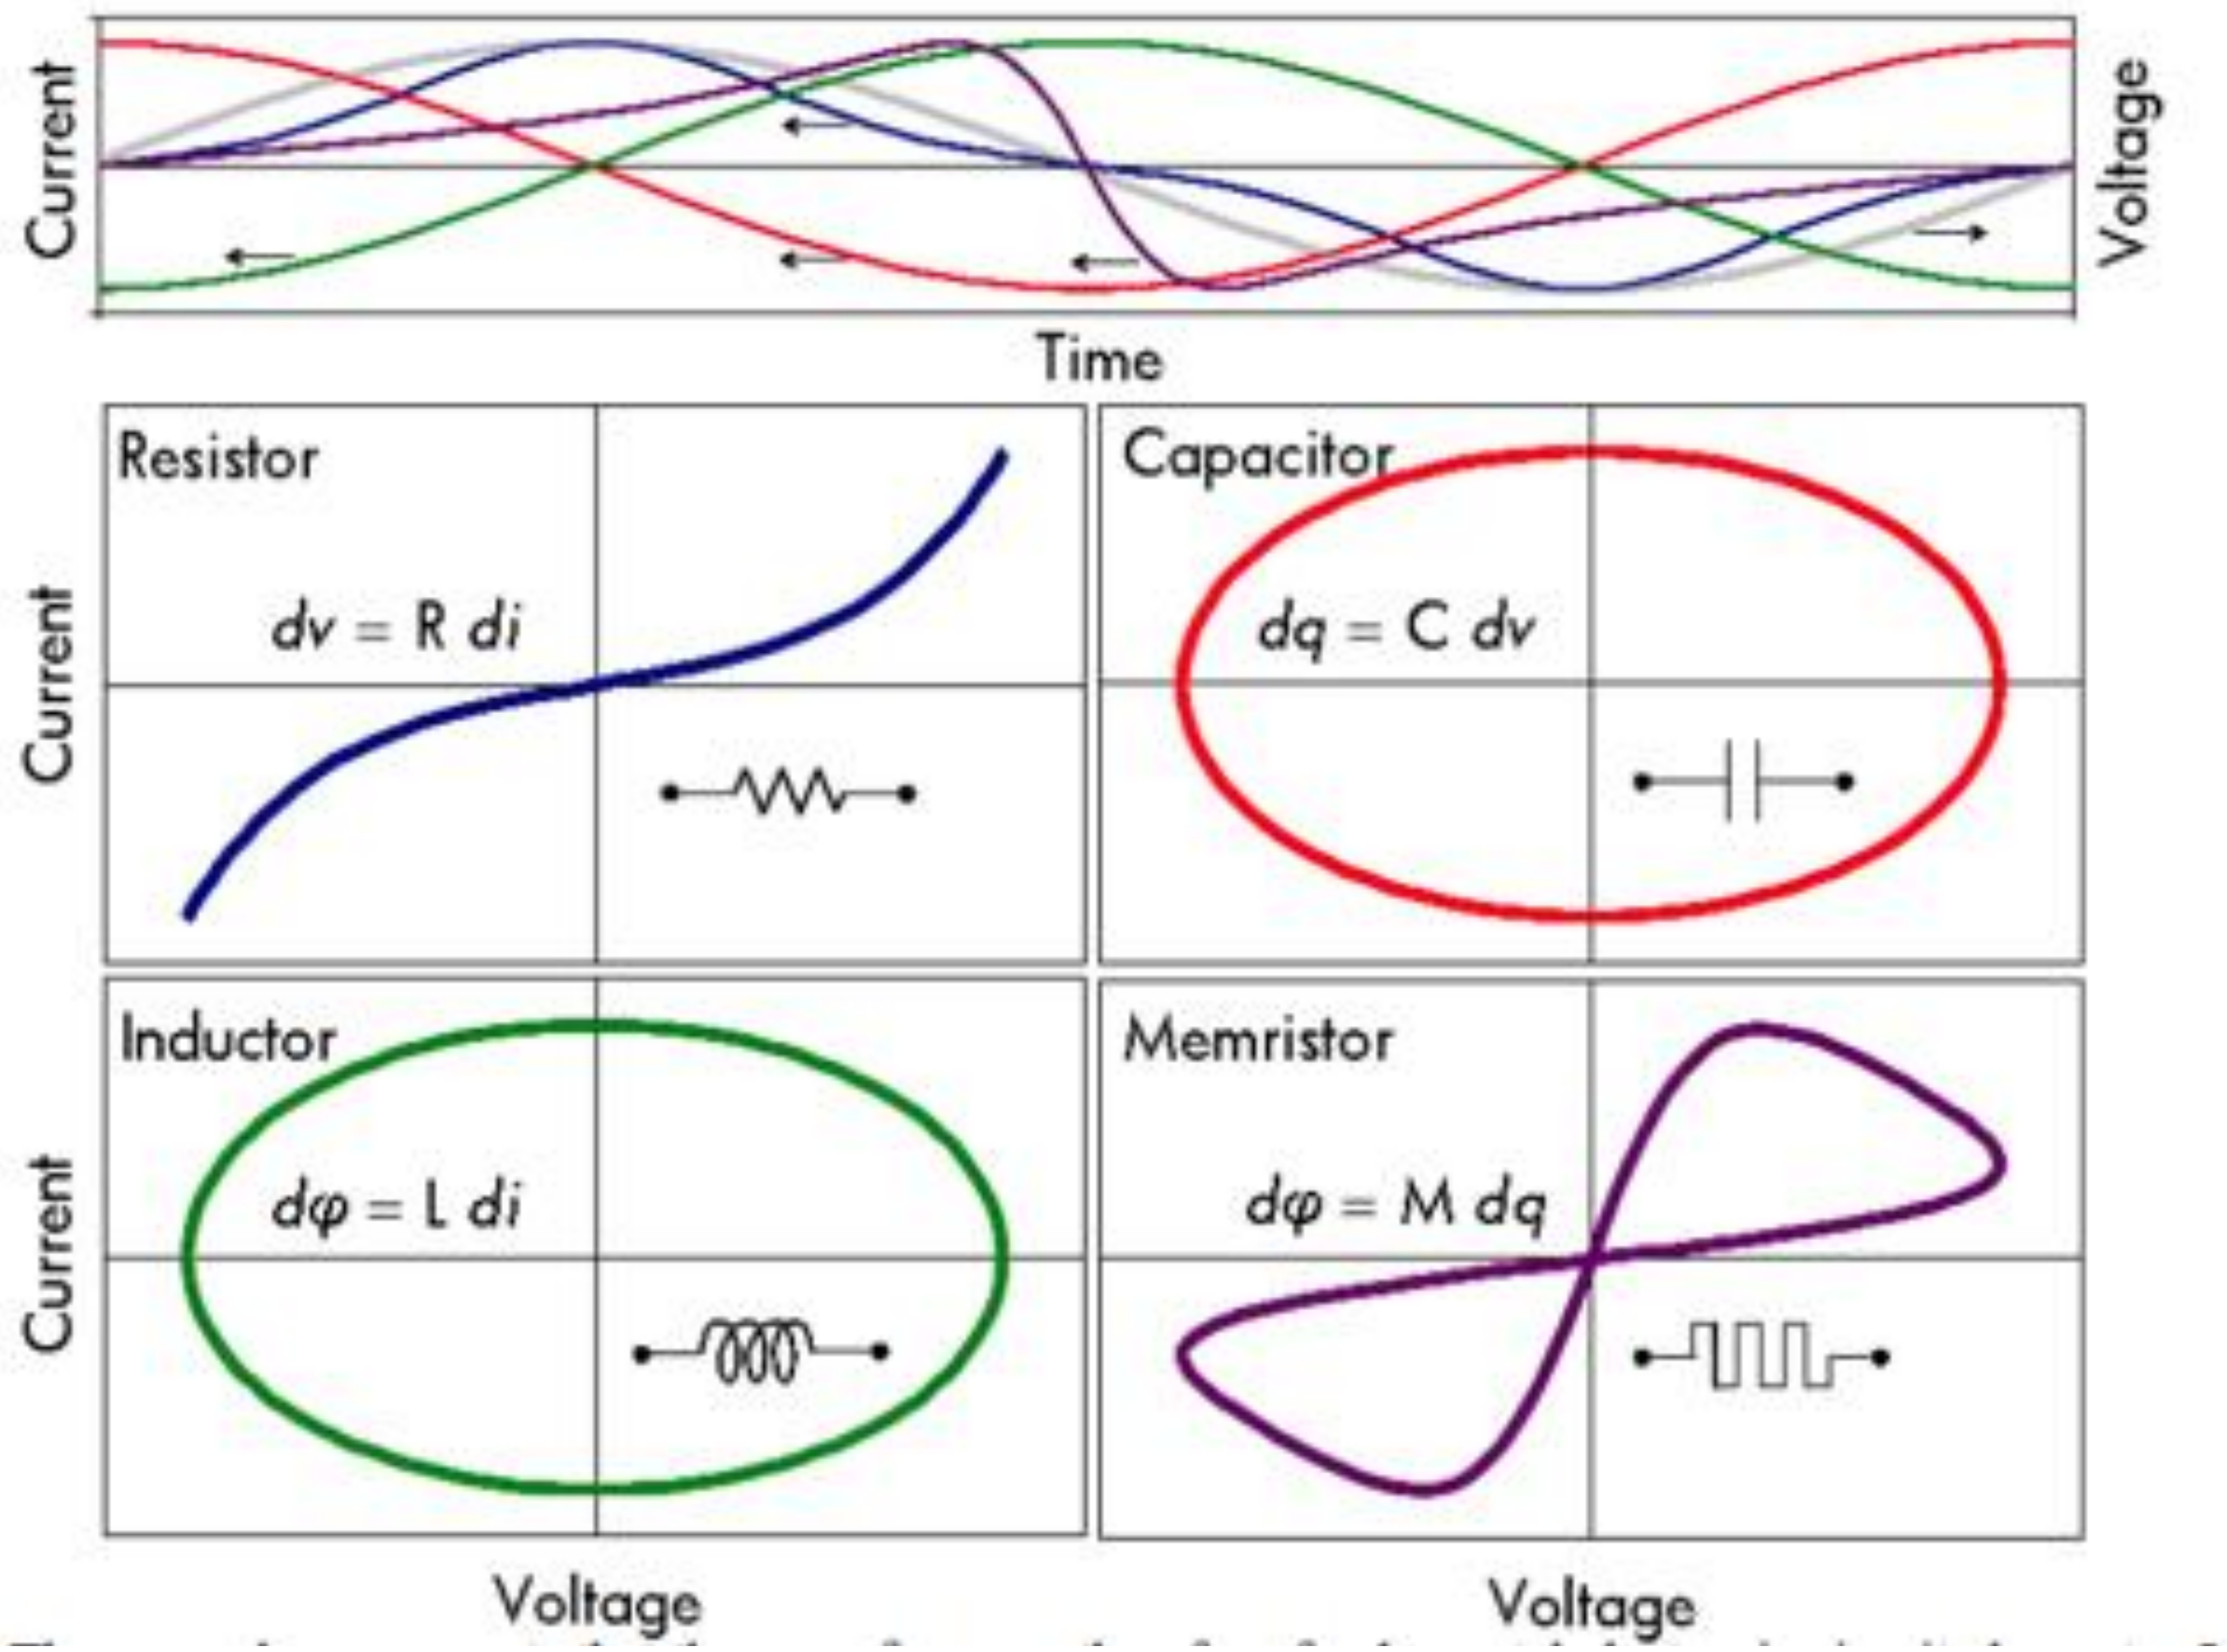

KNOWM

MemComputing, Inc.

The Good News

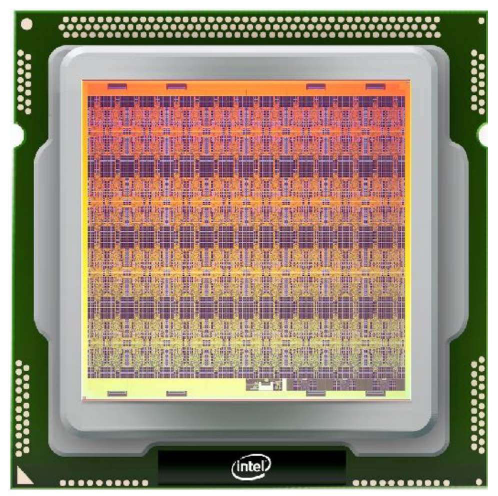

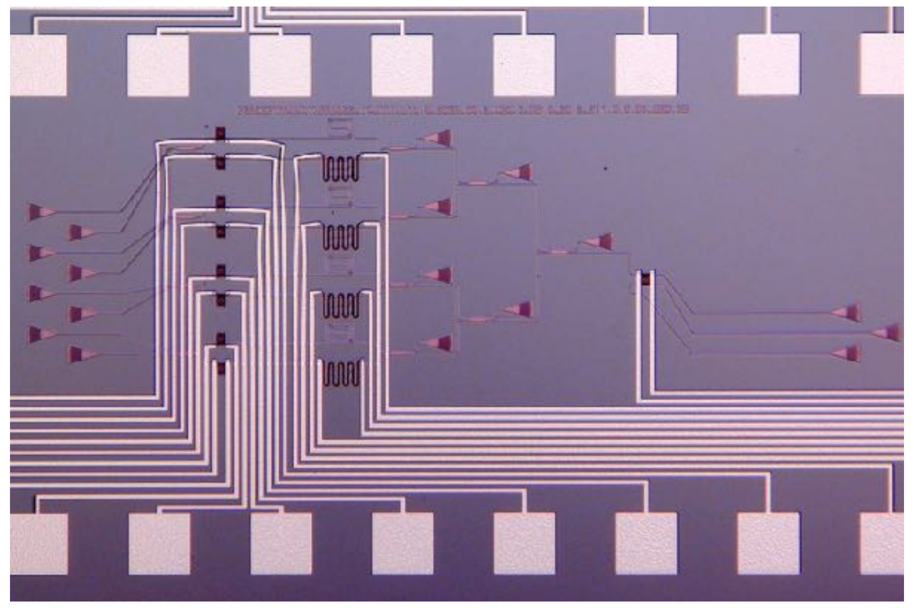

$$\sigma_i = 2x_i - 1$$

QUBO

$$\min : \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij} x_i x_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_i x_i$$

s.t.

$$x_i \in \{0, 1\} \;\forall i \in \mathcal{N}$$

Ising

$$\min : \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij} \sigma_i \sigma_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_i \sigma_i$$

s.t.

$$\sigma_i \in \{-1, 1\} \;\forall i \in \mathcal{N}$$

$$x_i = \frac{\sigma_i + 1}{2}$$

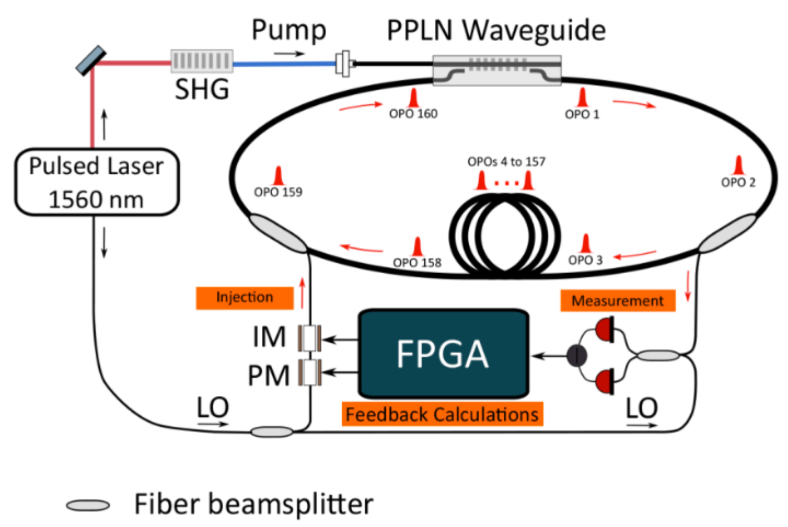

D:Wave

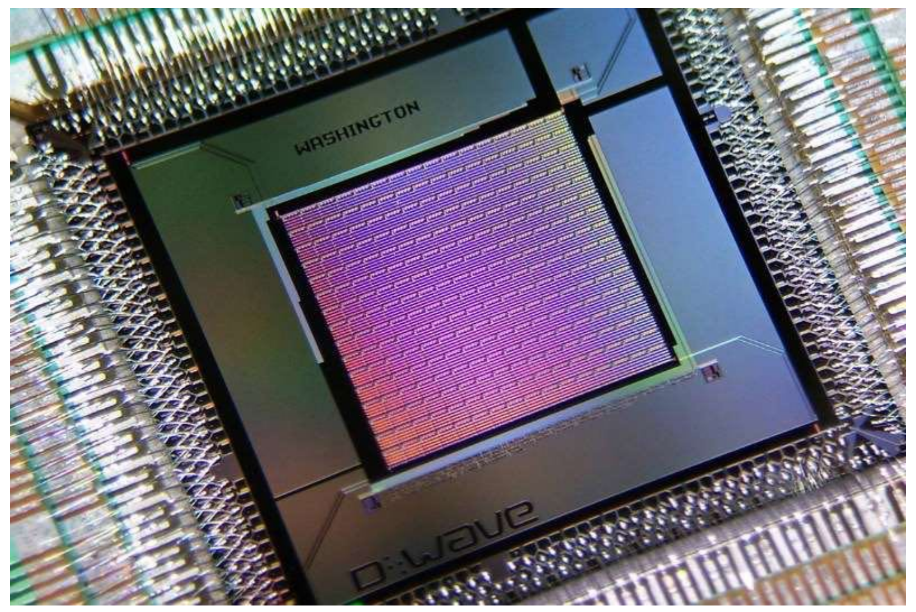

Ising Processing Units:
Potential and Challenges for Discrete Optimization

https://arxiv.org/pdf/1707.00355.pdf

The Ambition of Novel Computing

Computational Supremacy

"[Can a novel computing device] perform a well-defined computational task beyond the capabilities of state-of-the-art classical computers" Boixo et. al.

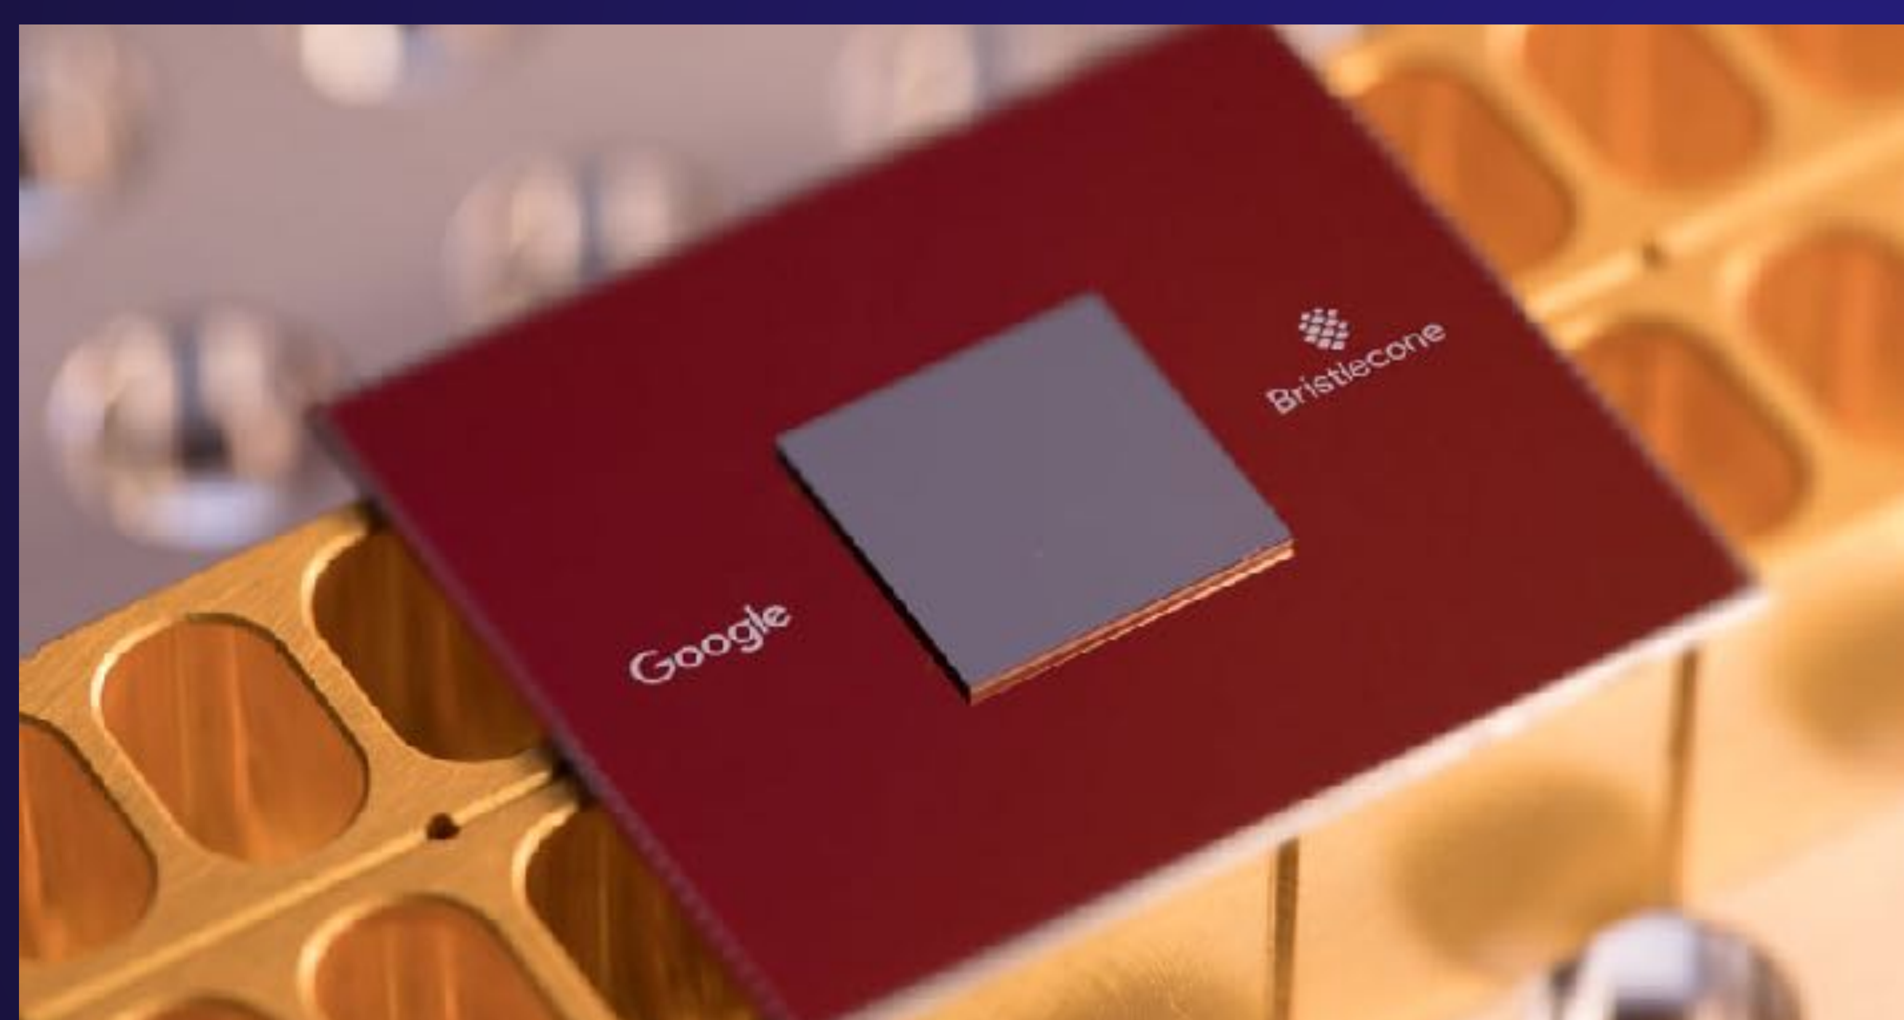

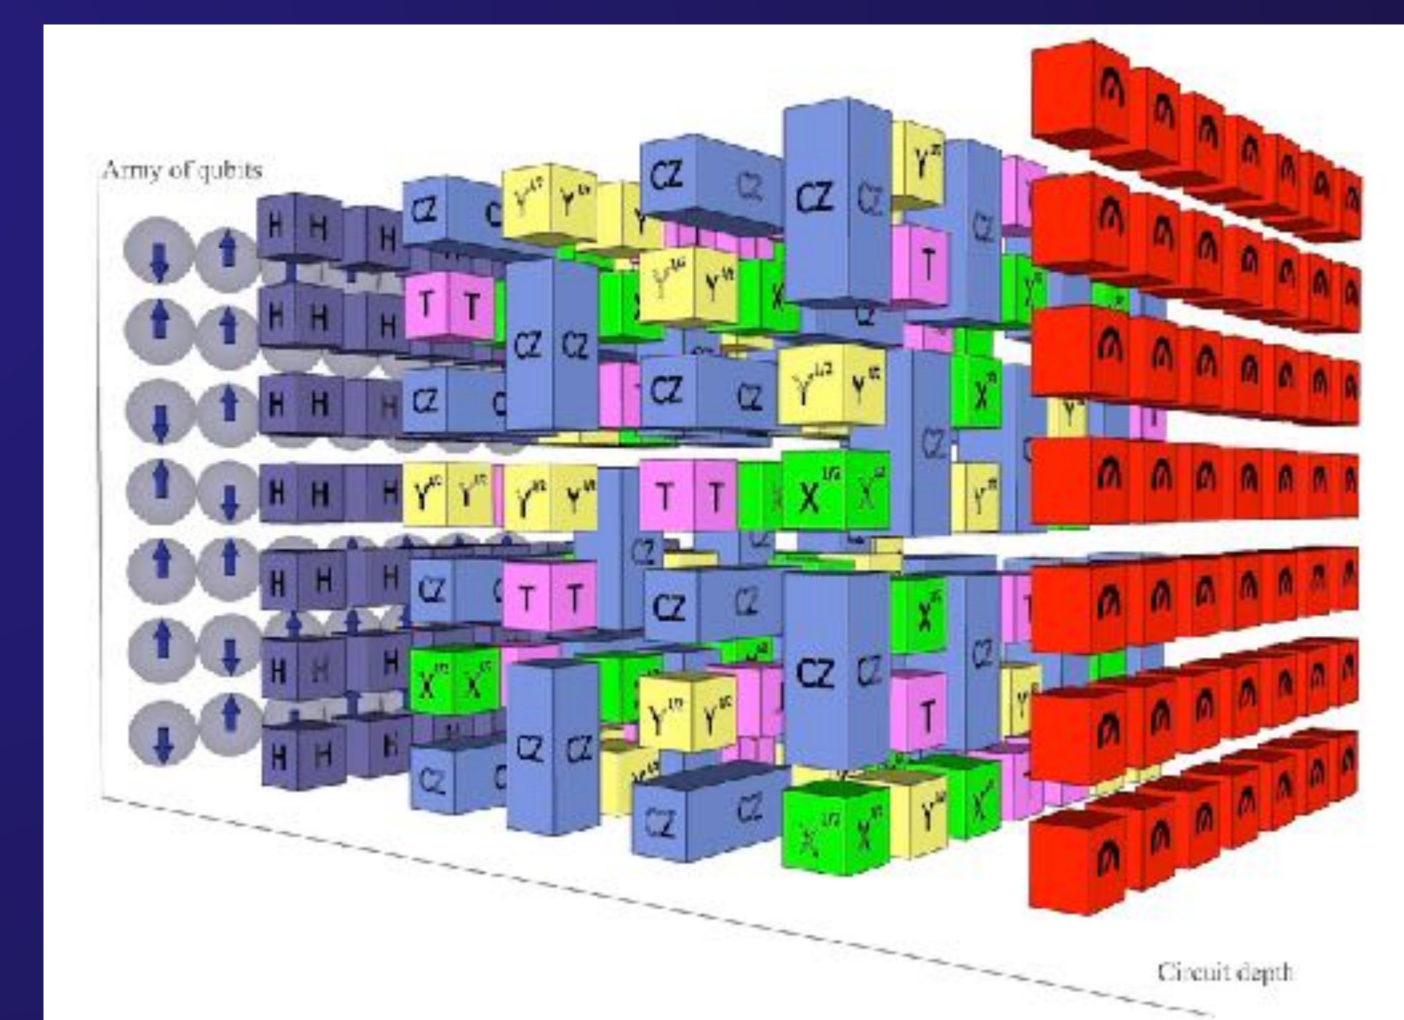

NOTE: Initial computational tasks are likely to be very contrived!

https://arxiv.org/pdf/1608.00263.pdf

https://ai.googleblog.com/2018/05/the-question-of-quantum-supremacy.html

The State of Novel Computing

UNIVAC 1960

The State of Novel Computing

What does it compute?

How fast does it compute?
(esp. compared to state-of-the-art)

Overview

- My experience trying to make sense of a D-Wave Quantum Annealer

- What does it compute?

- Benchmarking Successes and Failures

- Future Outlooks

Ising

$$\min: \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij} \sigma_i \sigma_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_i \sigma_i$$

s.t.

$$\sigma_i \in \{-1, 1\} \; \forall i \in \mathcal{N}$$

What does a D-Wave Compute?

The User Perspective

Ising Model Specification

$$\min : \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij}\sigma_i\sigma_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_i\sigma_i$$

s.t.

$$\sigma_i \in \{-1, 1\} \ \forall i \in \mathcal{N}$$

Runtime Parameters (e.g. # of replicates)

⇒ **D-Wave** ⇒

Variable Assignments

count	σ_1	σ_2	...	σ_n
342	-1	1	...	1
173	1	-1	...	-1
12	1	1	...	-1
...	...	...	...	...

Ground Truth…

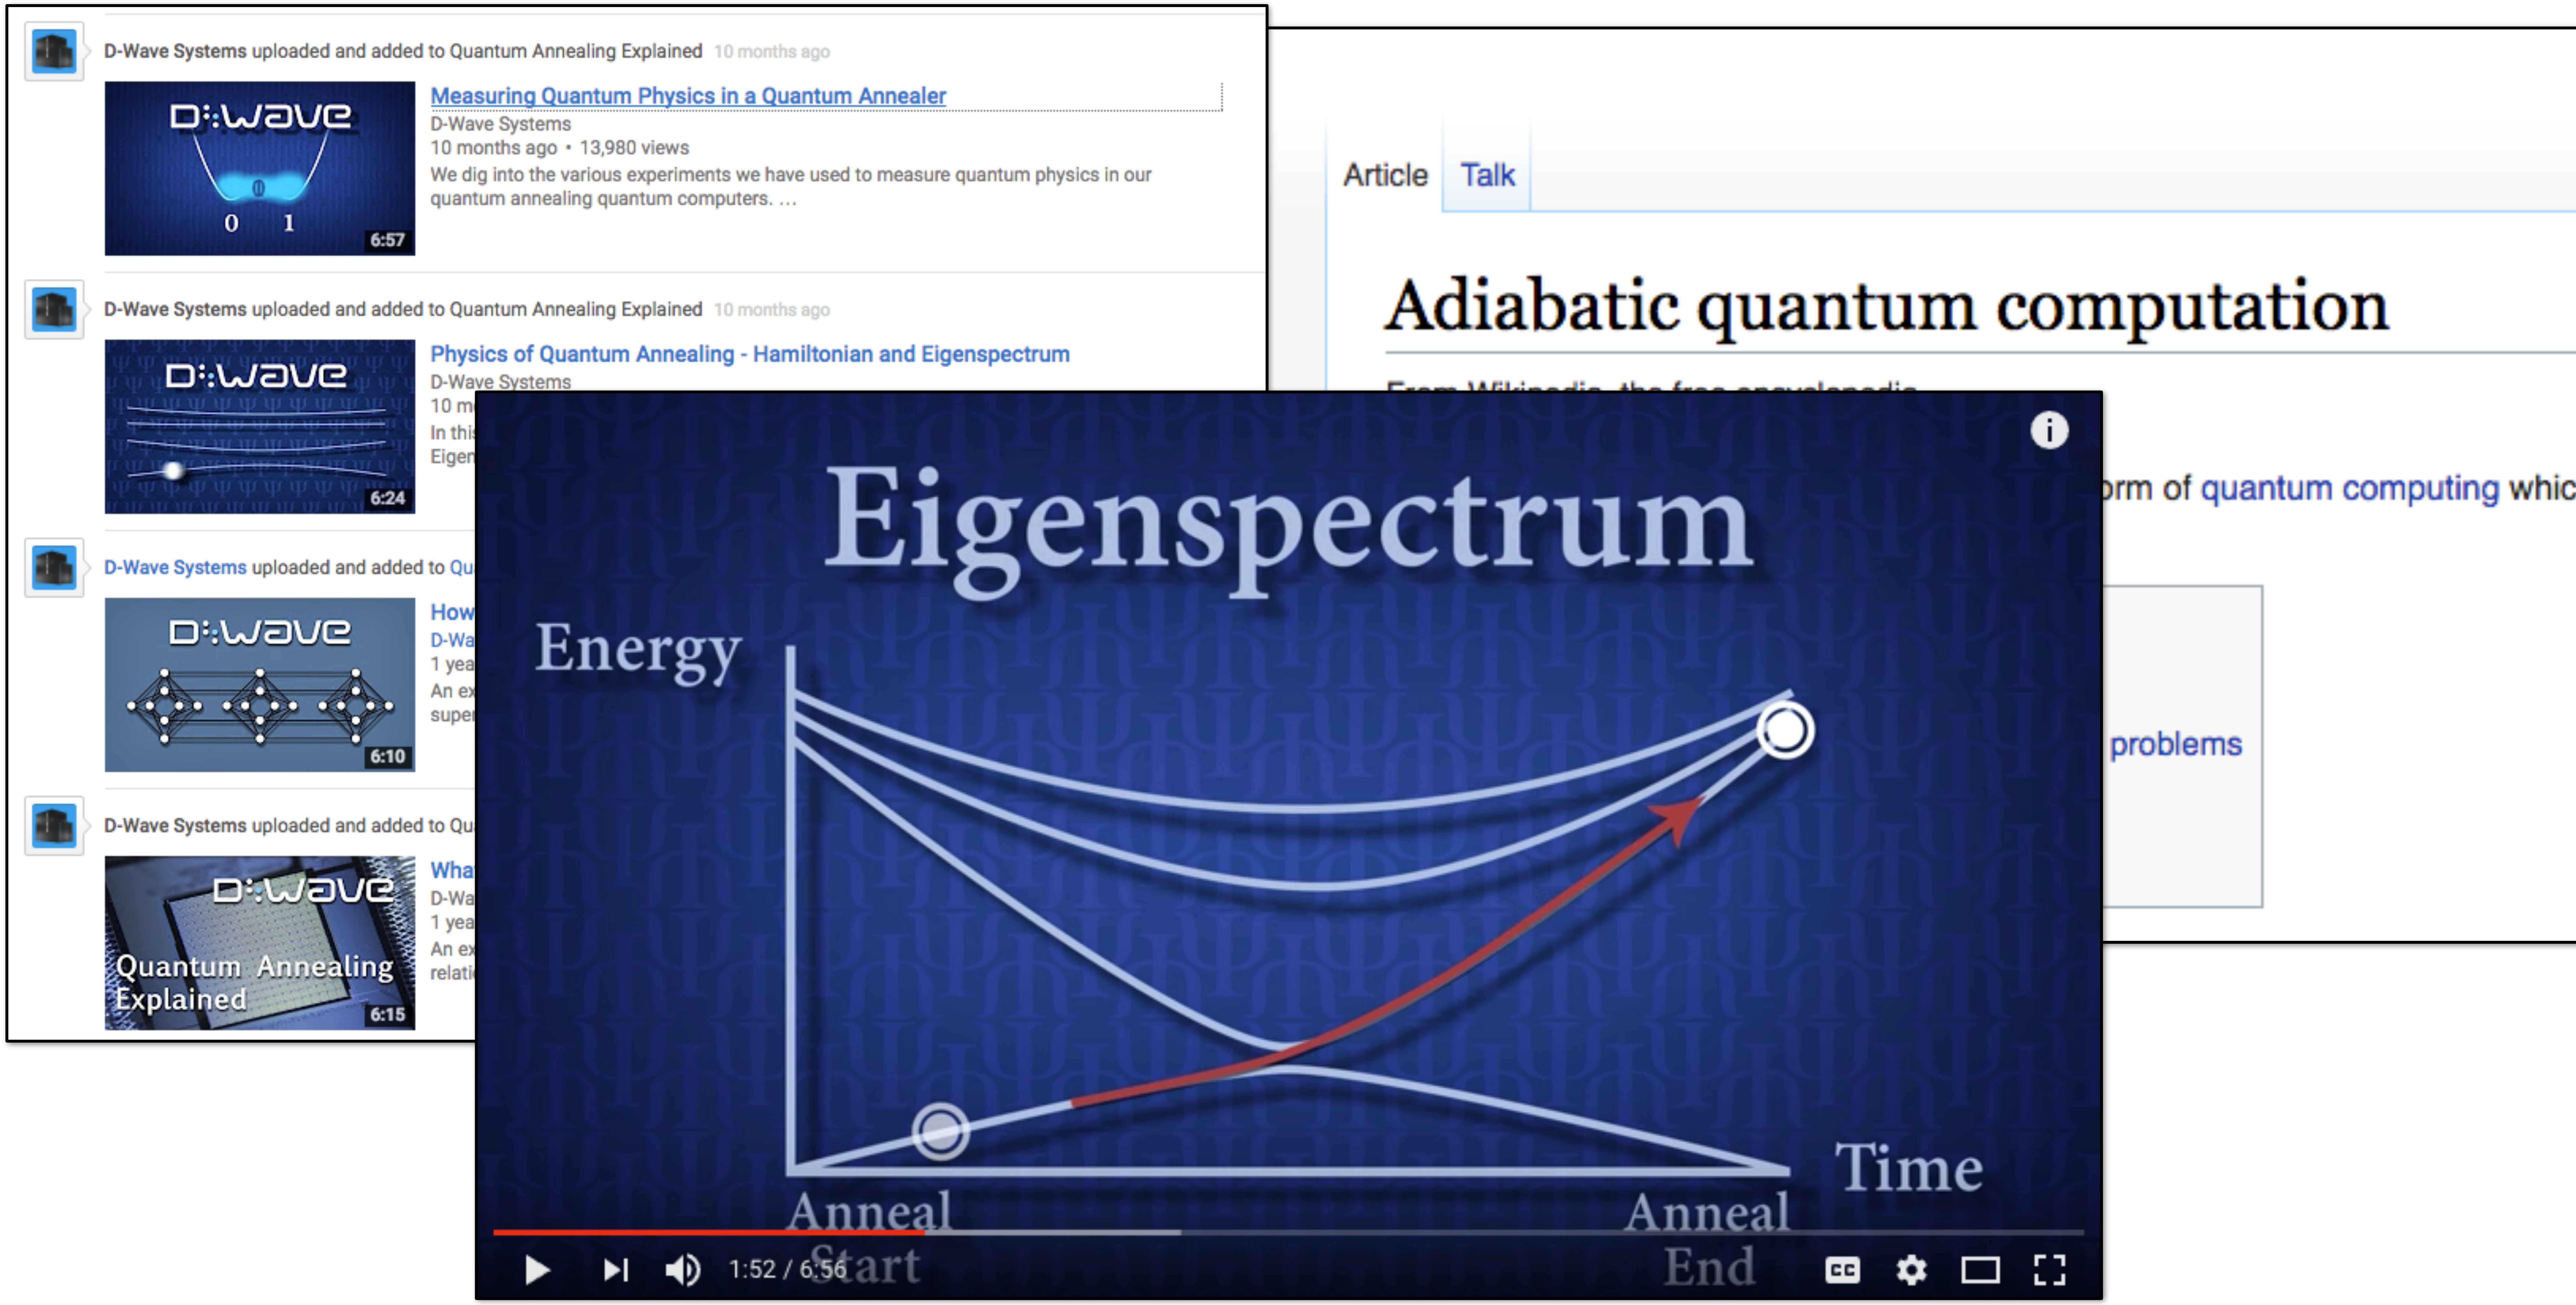

First Order Approximation

- Finds globally optimal solutions to the Ising Model,

So what?

$$\min : \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij}\sigma_i\sigma_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_i\sigma_i$$

s.t.

$$\sigma_i \in \{-1, 1\} \ \forall i \in \mathcal{N}$$

Max-Cut

$$\min : \sum_{i,j \in \mathcal{E}} \frac{\sigma_i\sigma_j - 1}{2}$$

s.t.

$$\sigma_i \in \{-1, 1\} \ \forall i \in \mathcal{N}$$

WOW!

Let’s Go!

Biq Mac Library - Binary quadratic and Max cut Library

This site offers a collection of Max-Cut instances and quadratic 0-1 programming problems of medium size. Most of the instances were collected while developing Biq Mac, an SDP based Branch & Bound code (see [RRW07] or [Wie06]). The dimension of the problems (i.e., number of variables or number of vertices in the graph) ranges from 20 to 500. The instances are mainly ment to be used for testing exact solution methods for quadratic 0-1 programming or Max-Cut problems.

Any comments or further instances to be added are welcome! Please contact angelika.wiegele@aau.at.

The structure of the directories is as follows:

```
                          |--- beasley
            |---- Biq ---|--- gka
            |             |--- be
            |
Biq Mac  ---|
            |             |--- rudy
            |---- Mac ---|
                          |--- ising
```

Some challenges…

http://biqmac.uni-klu.ac.at/biqmaclib.html

Catch #1: Hardware Graph

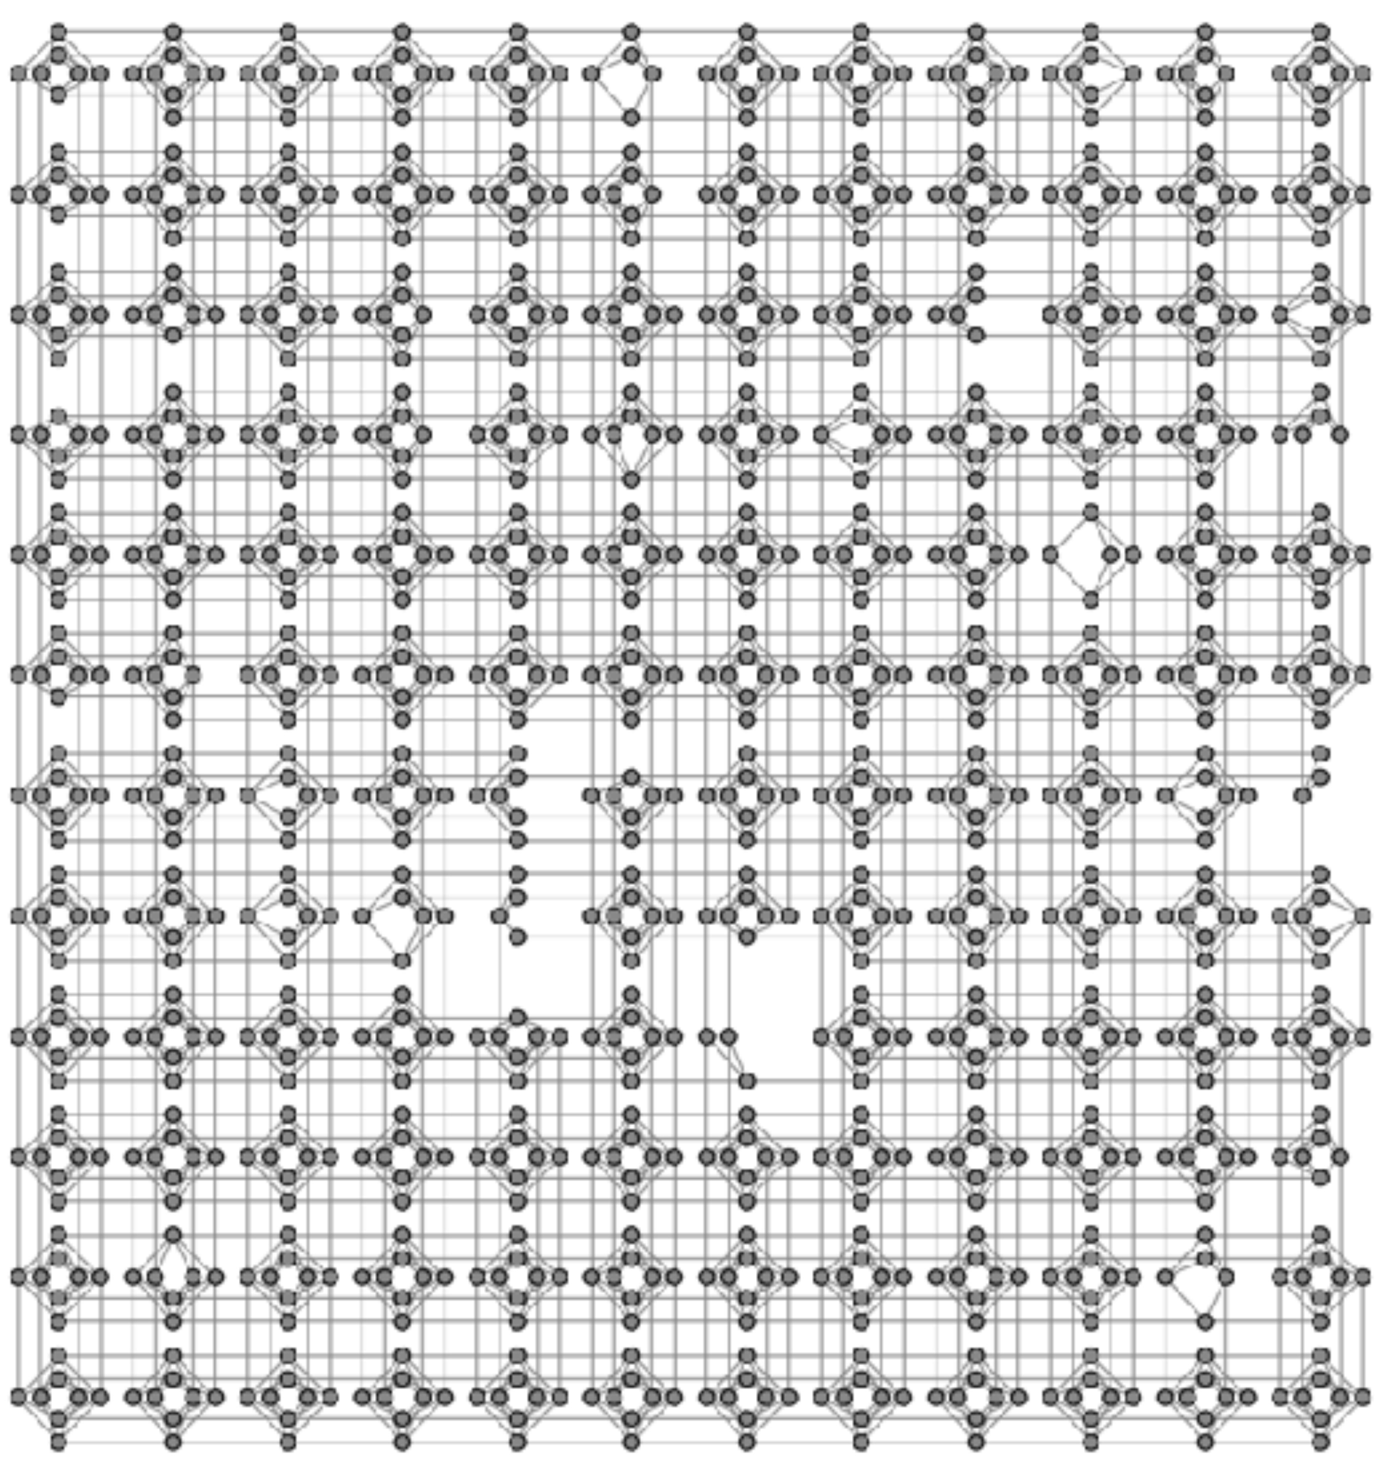

$\mathcal{C}_{12}$ DW2X

$$\min : \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij} \sigma_i \sigma_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_i \sigma_i$$

s.t.

$$\sigma_i \in \{-1, 1\} \; \forall i \in \mathcal{N}$$

$$\mathcal{E} \subseteq \mathcal{C}_{12}$$

$$|\mathcal{N}| \leq 1100$$

$$|\mathcal{E}| \leq 3068$$

Catch #2: Coefficient Values

$$\boldsymbol{c}_{ij} \in \{-1.00, -0.99, -0.98, \ldots, 0.98, 0.99, 1.00\} \; \forall i, j \in \mathcal{E}$$
$$\boldsymbol{c}_{i} \in \{-2.00, -1.99, -1.98, \ldots, 1.98, 1.99, 2.00\} \; \forall i \in \mathcal{N}$$

$$\min : \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij} \sigma_i \sigma_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_i \sigma_i$$

s.t.

$$\sigma_i \in \{-1, 1\} \; \forall i \in \mathcal{N}$$
$$\mathcal{E} \subseteq \mathcal{C}_{12}$$

Catch #3: Graph Embedding

Source Graph

qblib_3867

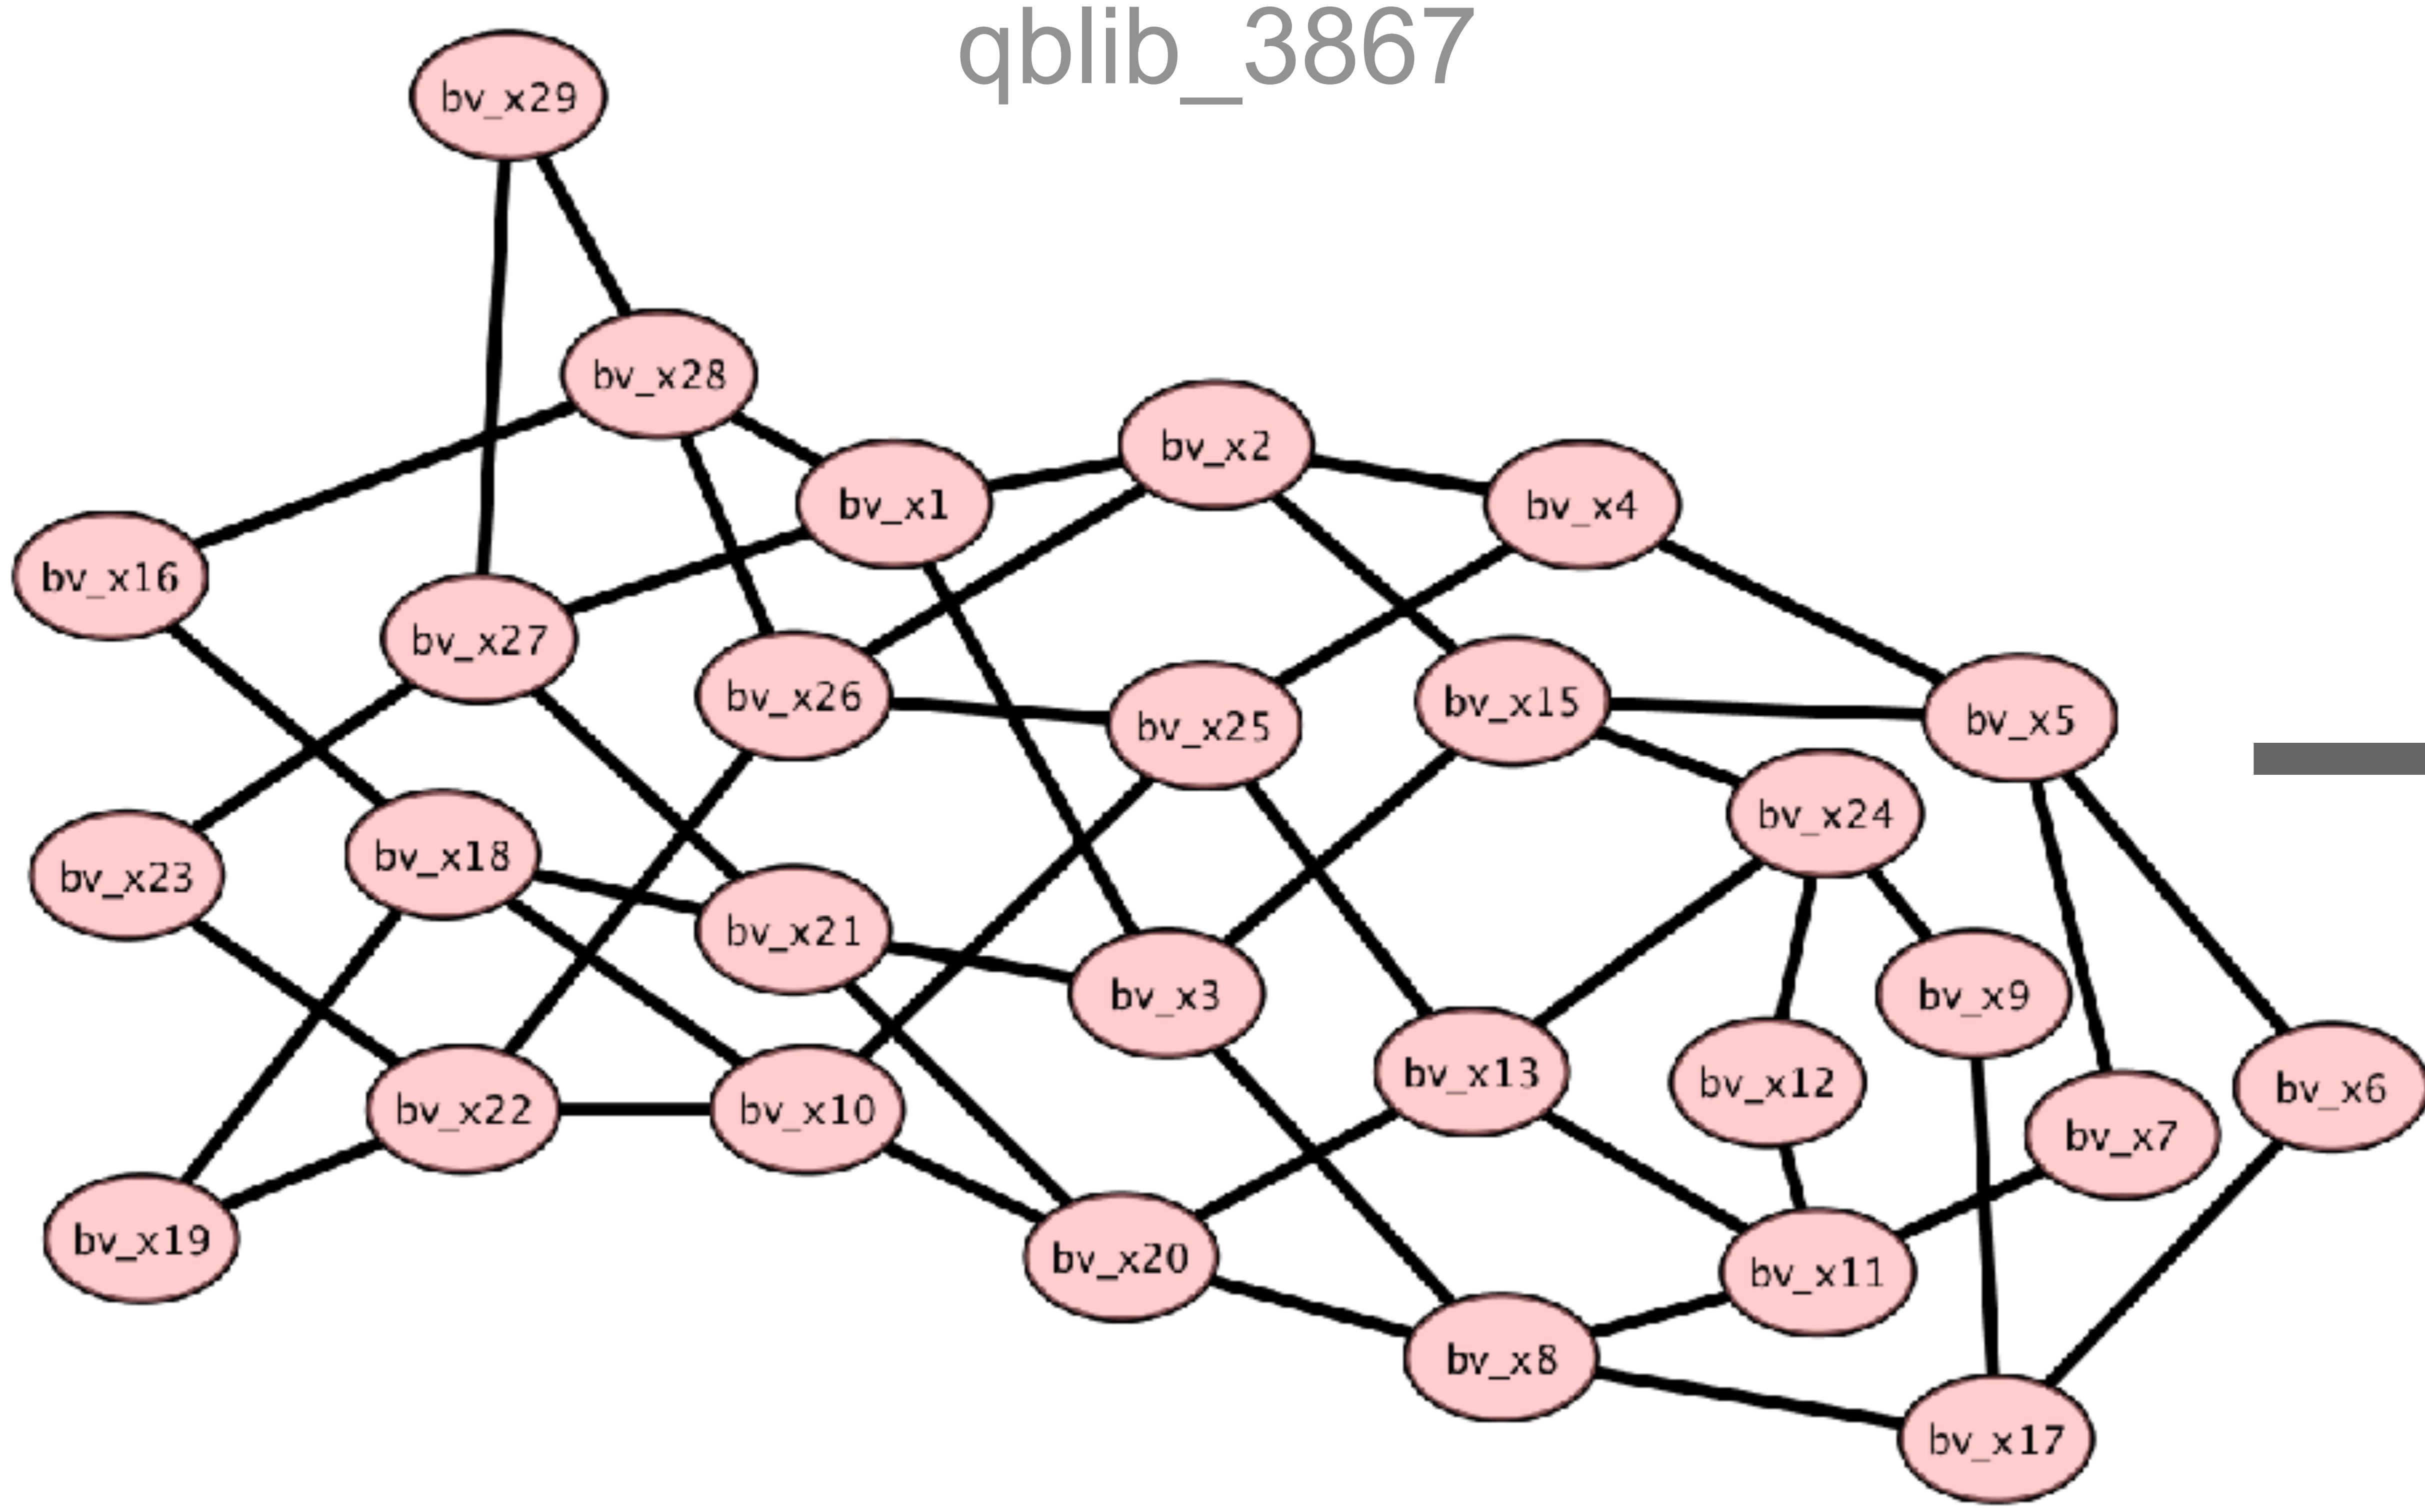

Target Graph

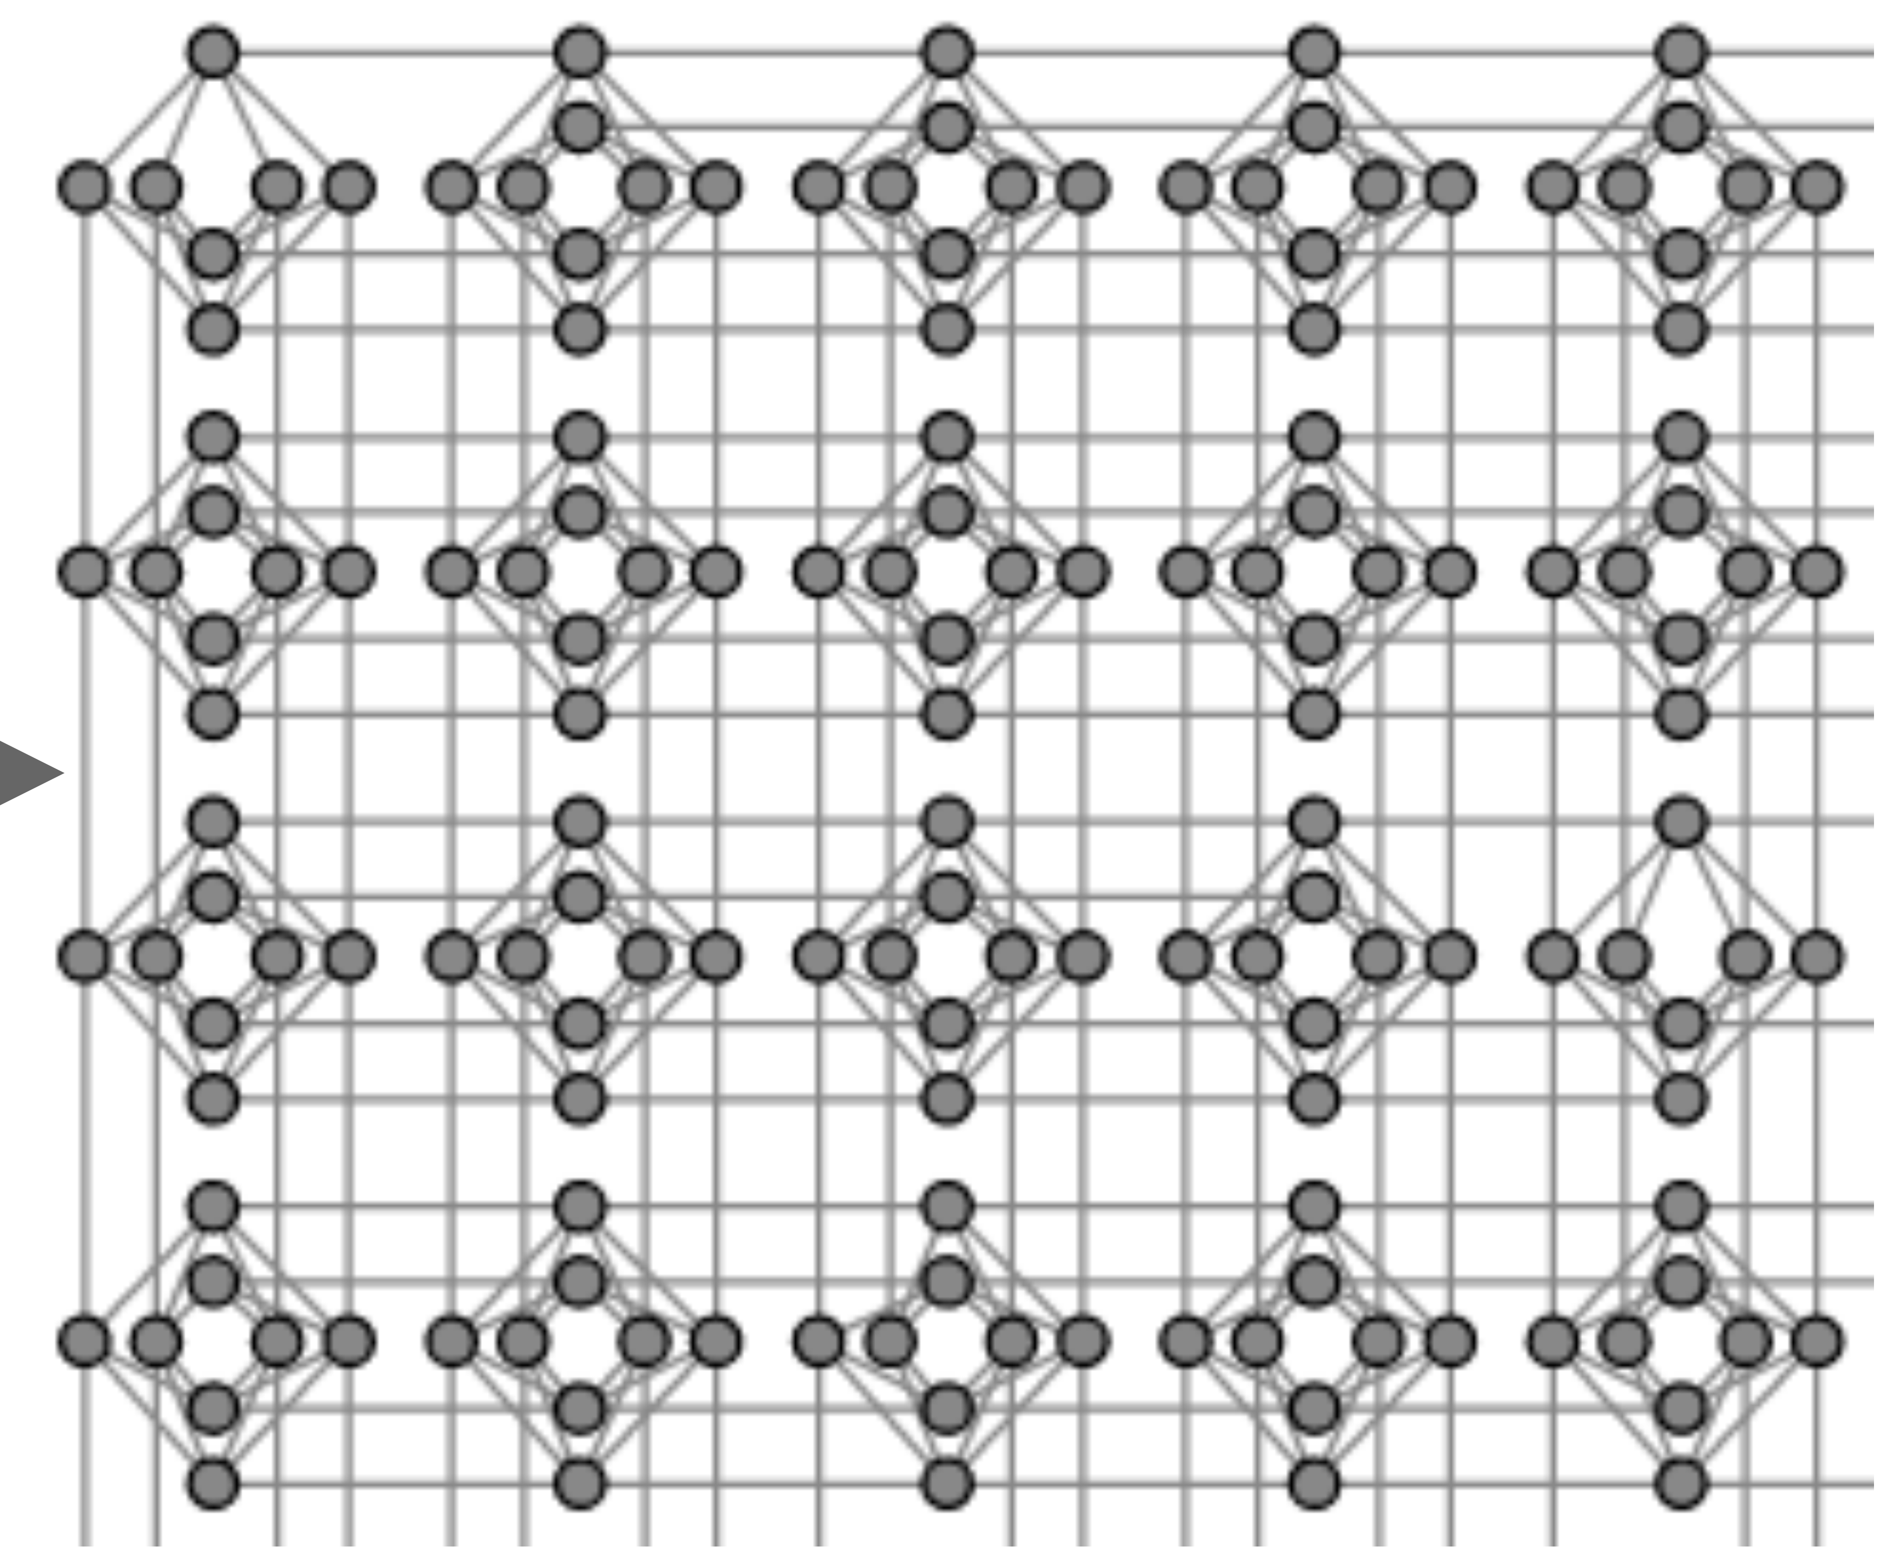

Catch #3: Graph Embedding

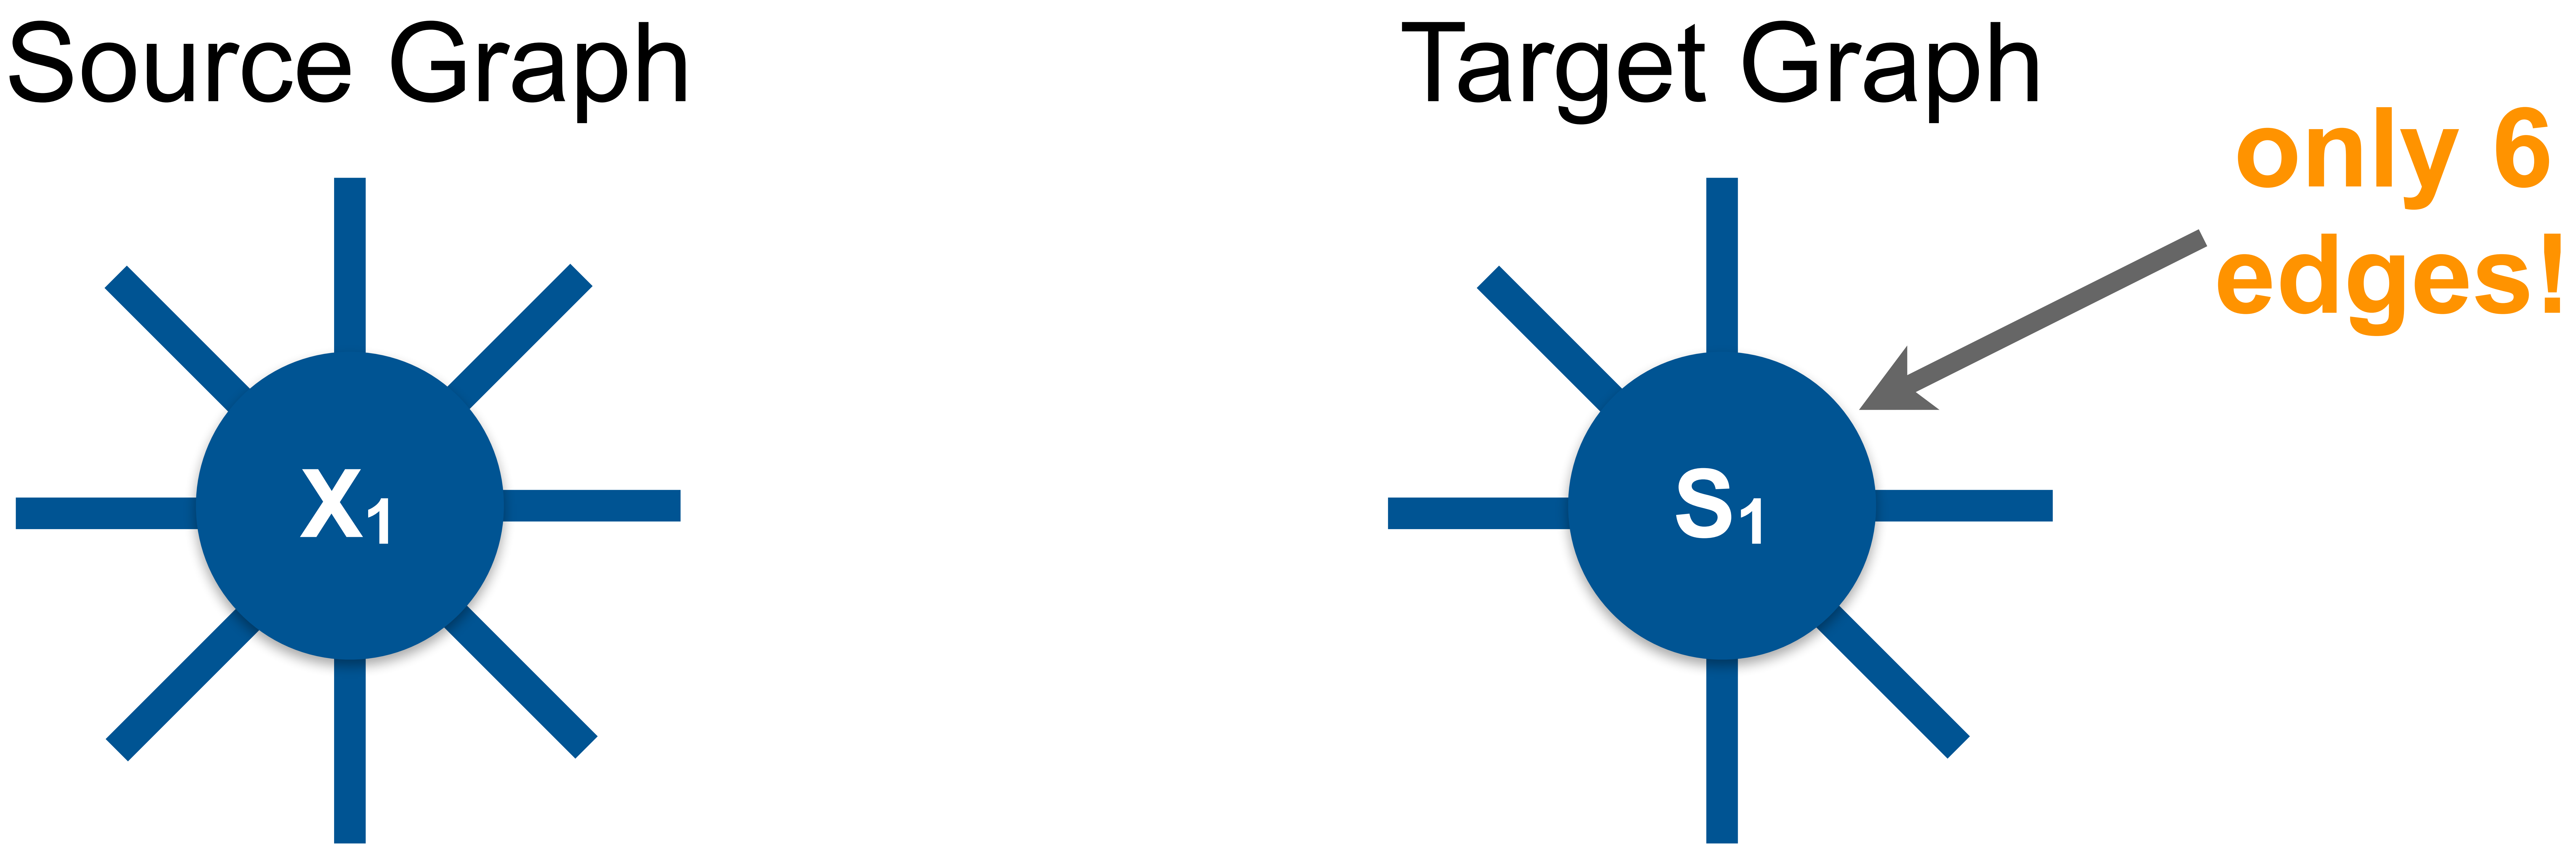

Catch #3: Graph Embedding

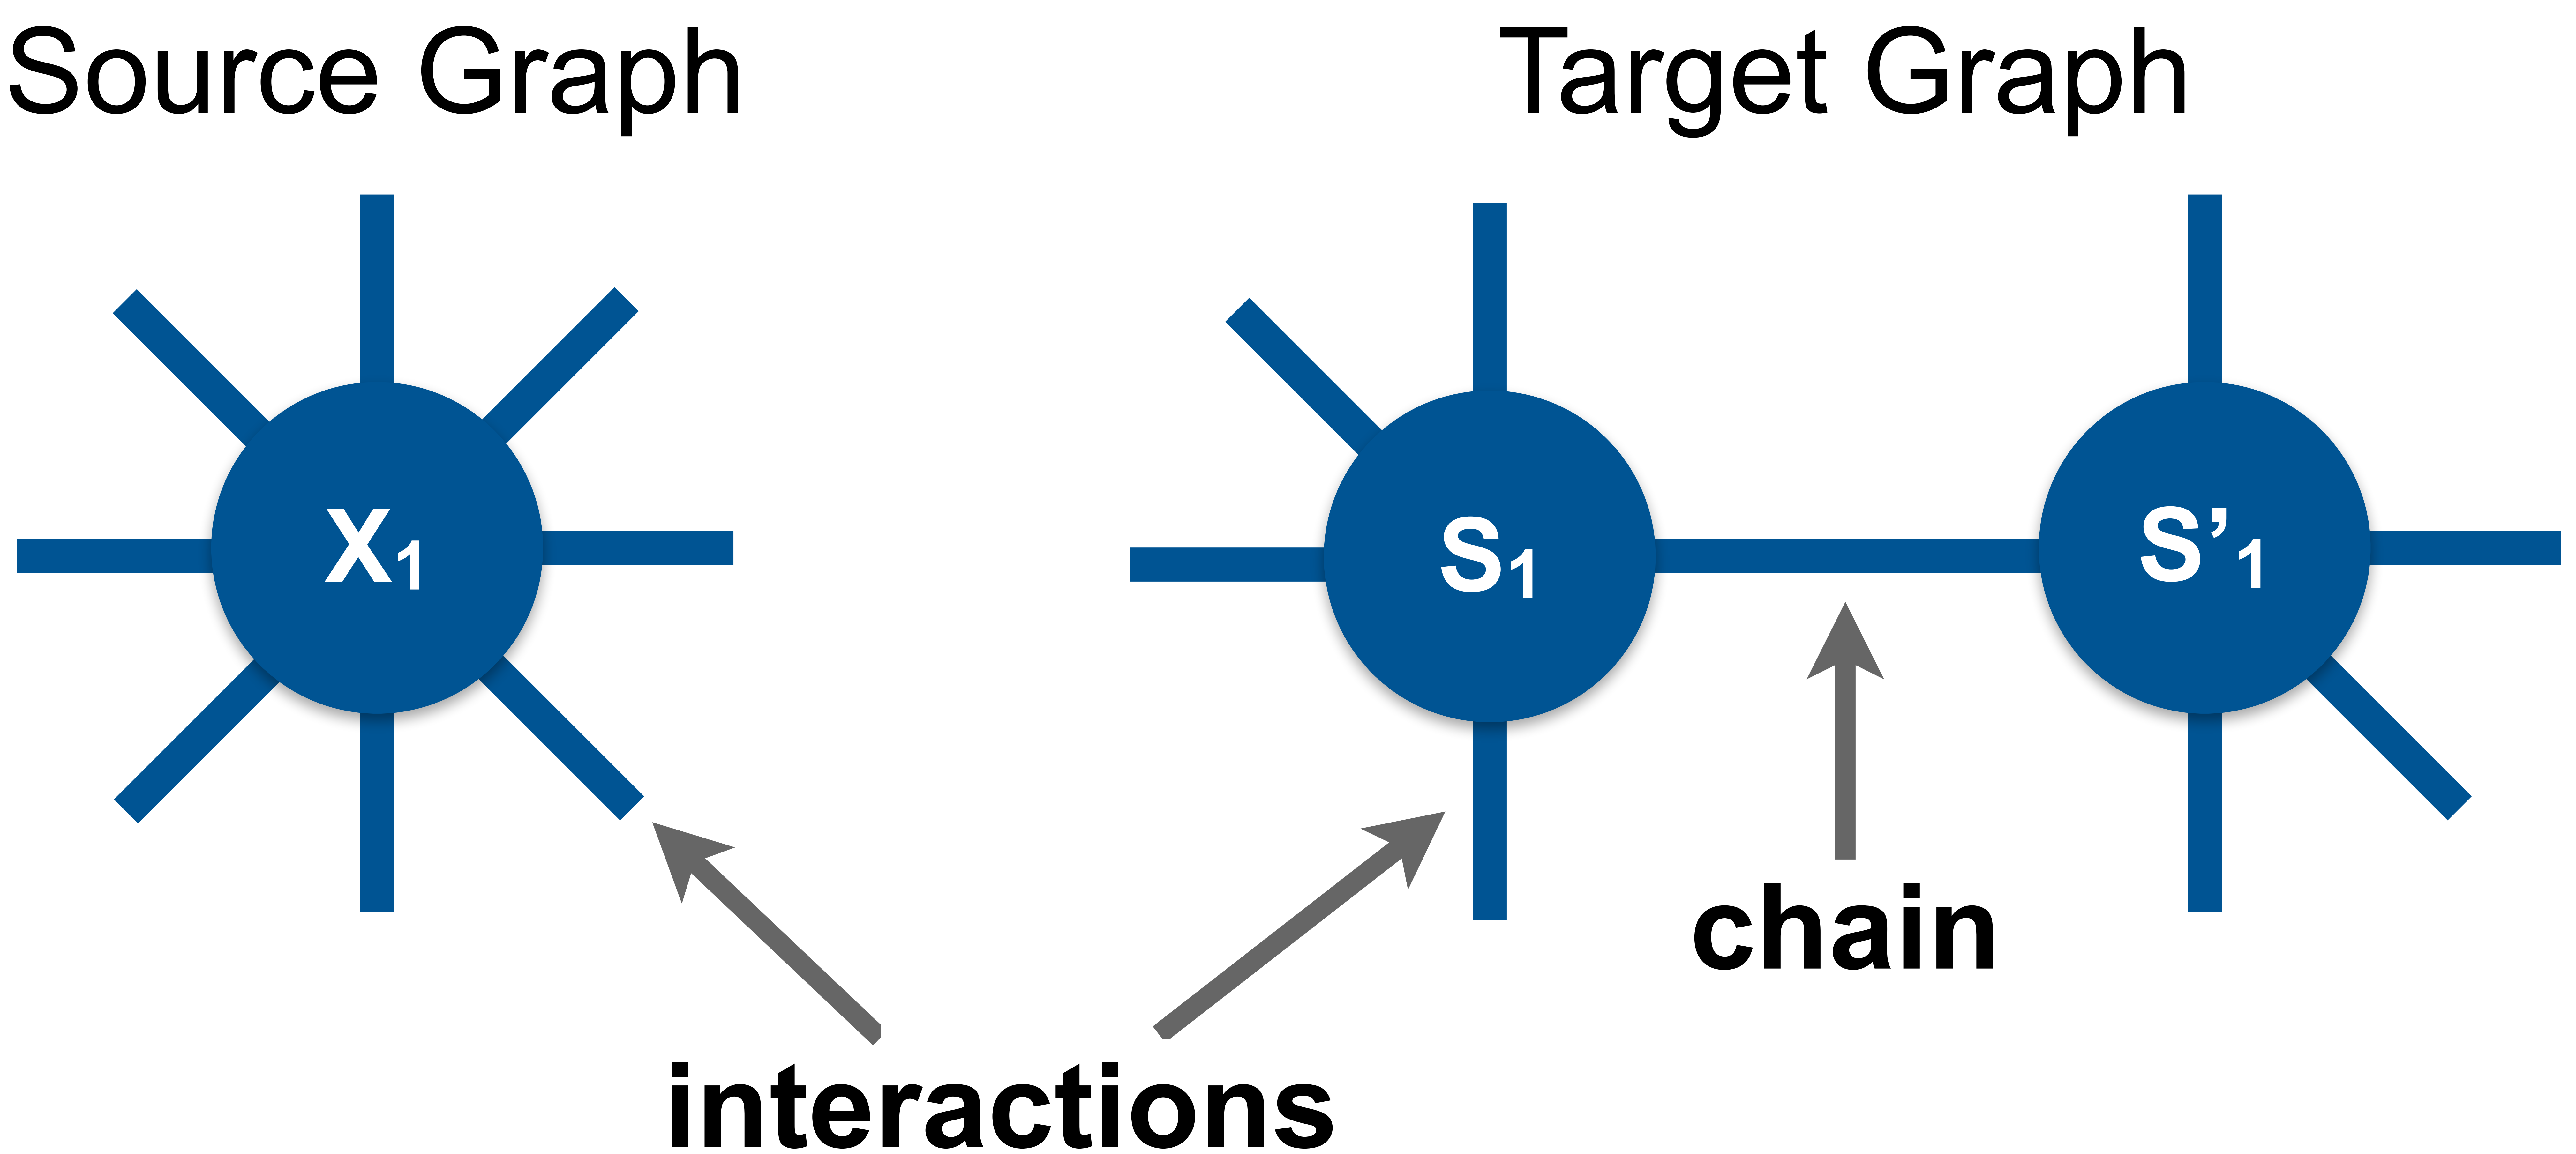

Warning: broken chains = infeasible solution

Typical D-Wave Algorithm

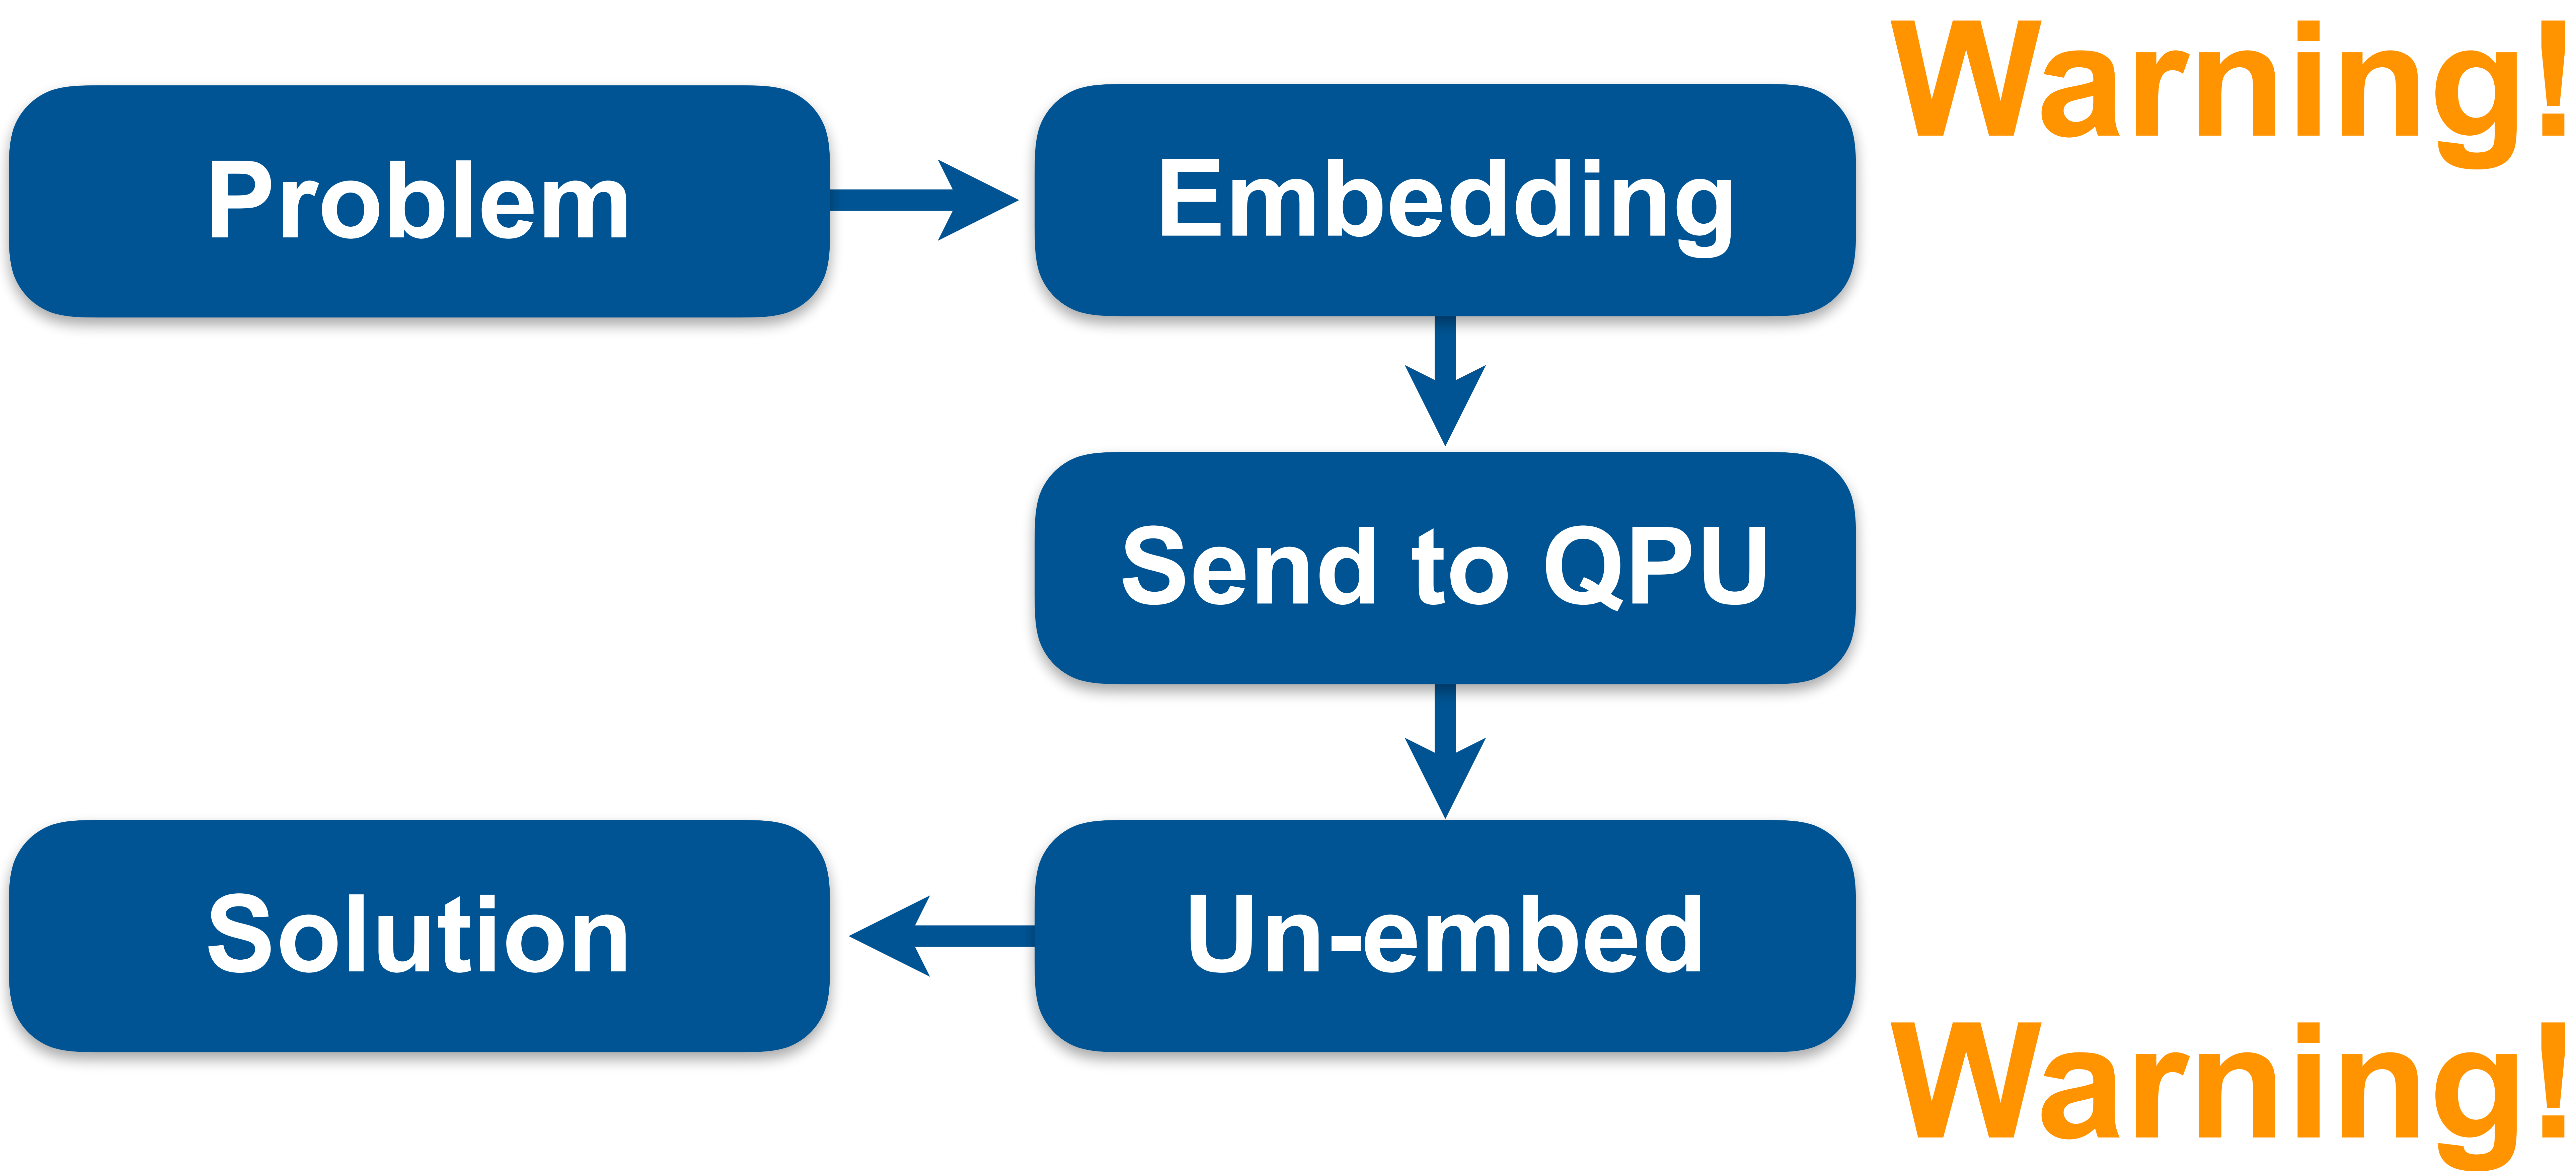

A Note on Other Technologies

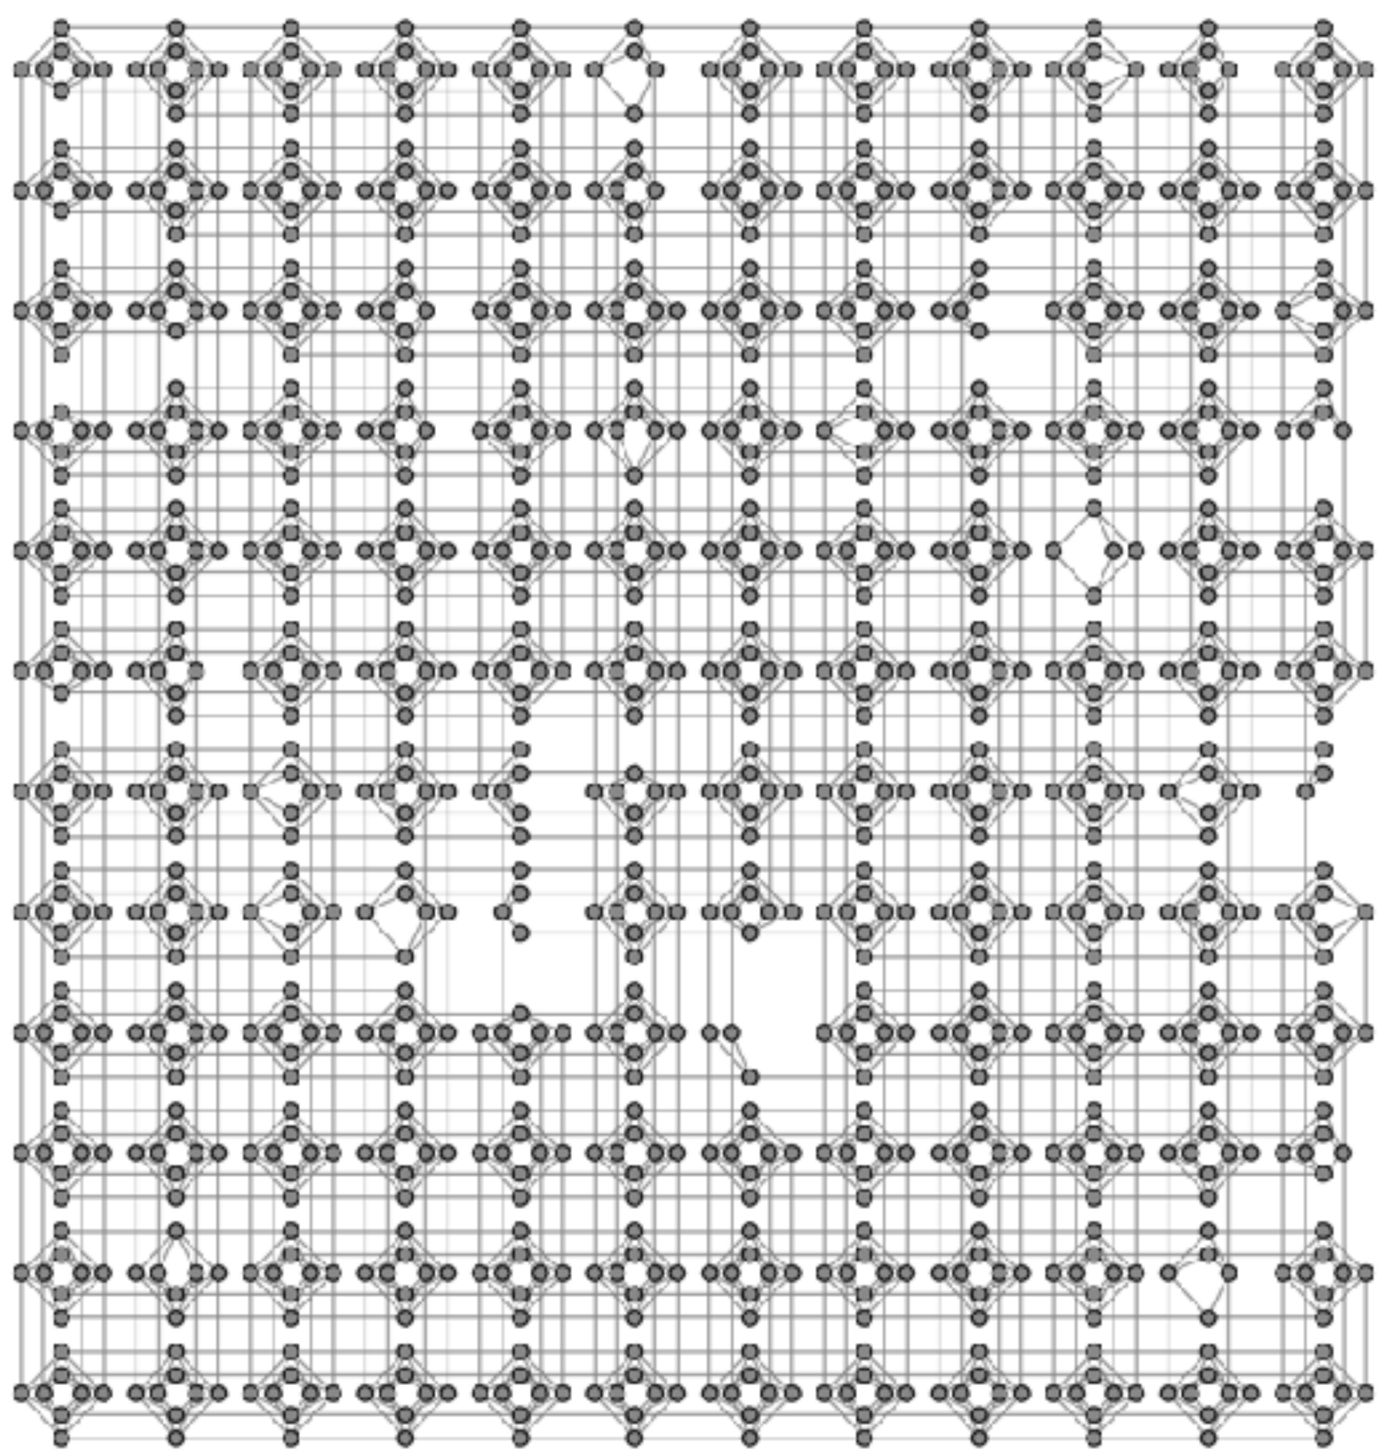

$$\boldsymbol{c}_{ij} \in \{-1.00, -0.99, -0.98, \ldots, 0.98, 0.99, 1.00\}\ \forall i, j \in \mathcal{E}$$
$$\boldsymbol{c}_{i} \in \{-2.00, -1.99, -1.98, \ldots, 1.98, 1.99, 2.00\}\ \forall i \in \mathcal{N}$$

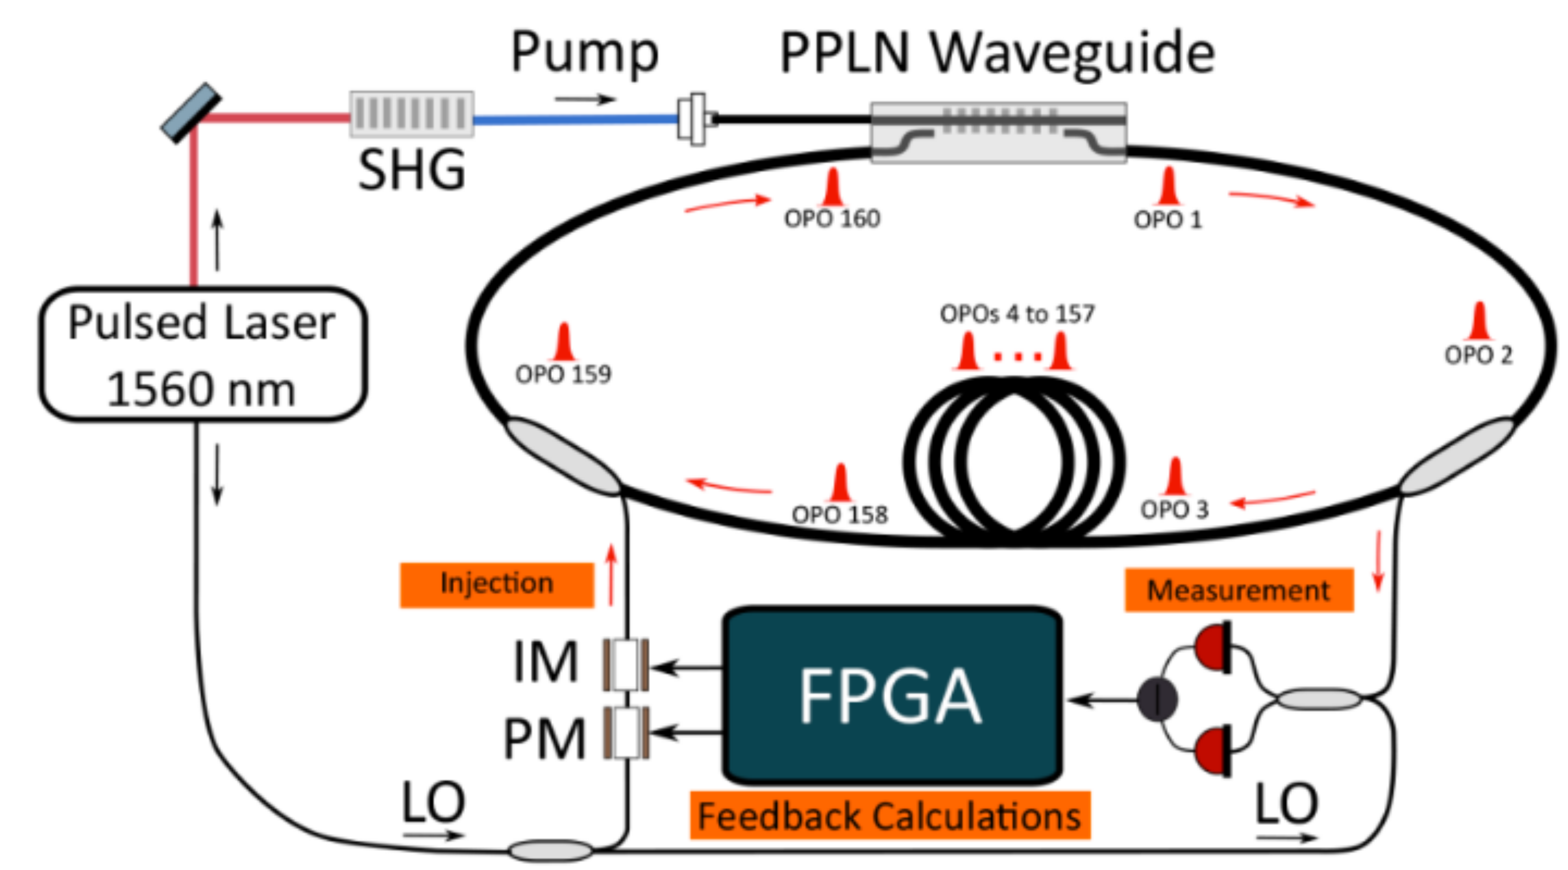

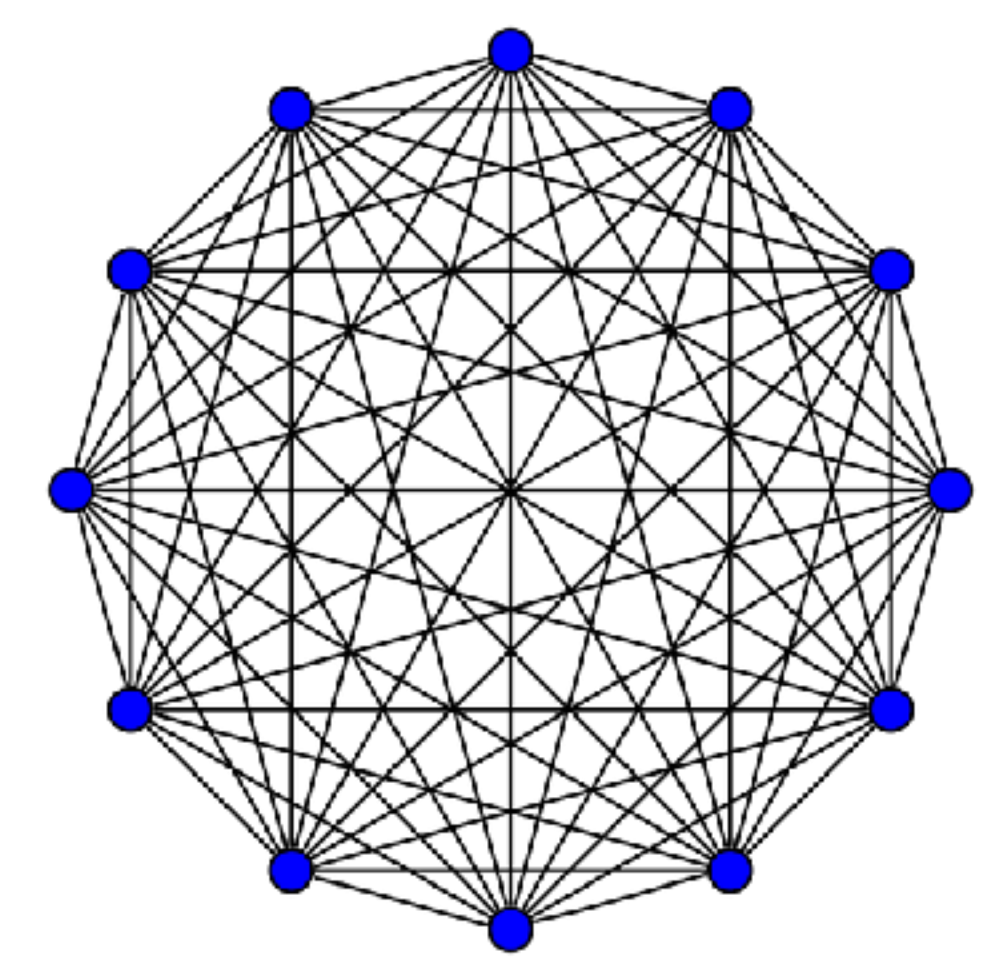

$$\boldsymbol{c}_{ij} \in \{-1, 0, 1\}\ \forall i, j \in \mathcal{E}$$
$$\boldsymbol{c}_{i} \in \{-1, 0, 1\}\ \forall i \in \mathcal{N}$$

Back to Benchmarking

Classic Operations Research Approach

- **Choose a widely recognized benchmark library**
  - QPLib
  - DIMACS Max-Clique Cases

QPLIB

- **Measure runtime against state-of-the-art alternatives**
  - Complete Search (e.g. MIP)
  - Local Search Heuristics (e.g. HFS)

GUROBI
OPTIMIZATION

HFS

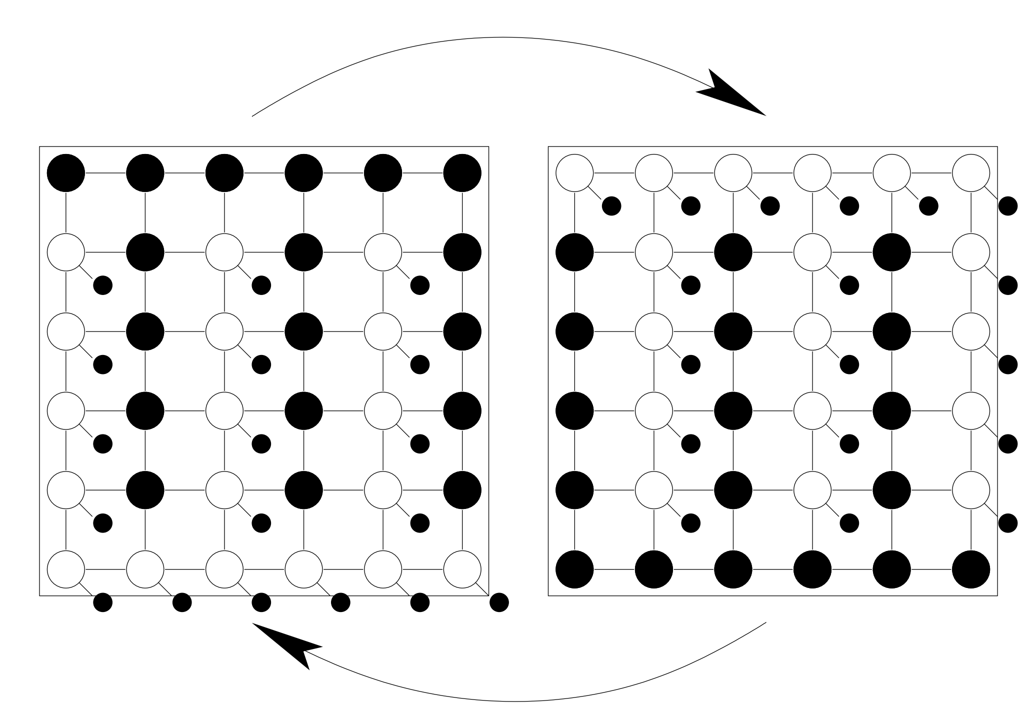

MIP-Based Ising Model Solver

Ising Model

$$\min : \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij}\sigma_i\sigma_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_i\sigma_i$$

s.t.

$$\sigma_i \in \{-1, 1\}\ \forall i \in \mathcal{N}$$

$$\mathcal{E} \subseteq \mathcal{C}_{12}$$

QUBO Model

$$\min : \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij}x_ix_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_ix_i$$

s.t.

$$x_i \in \{0, 1\}\ \forall i \in \mathcal{N}$$

$$\mathcal{E} \subseteq \mathcal{C}_{12}$$

ILP Model

$$\min : \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij}x_{ij} + \sum_{i \in \mathcal{N}} \boldsymbol{c}_ix_i$$

s.t.

$$2x_{ij} \geq x_i + x_j\ \forall i, j \in \mathcal{E}$$

$$x_{ij} \leq x_i\ \forall i, j \in \mathcal{E}$$

$$x_{ij} \leq x_j\ \forall i, j \in \mathcal{E}$$

$$x_{ij} \in \{0, 1\}\ \forall i, j \in \mathcal{E}$$

$$x_i \in \{0, 1\}\ \forall i \in \mathcal{N}$$

$$\mathcal{E} \subseteq \mathcal{C}_{12}$$

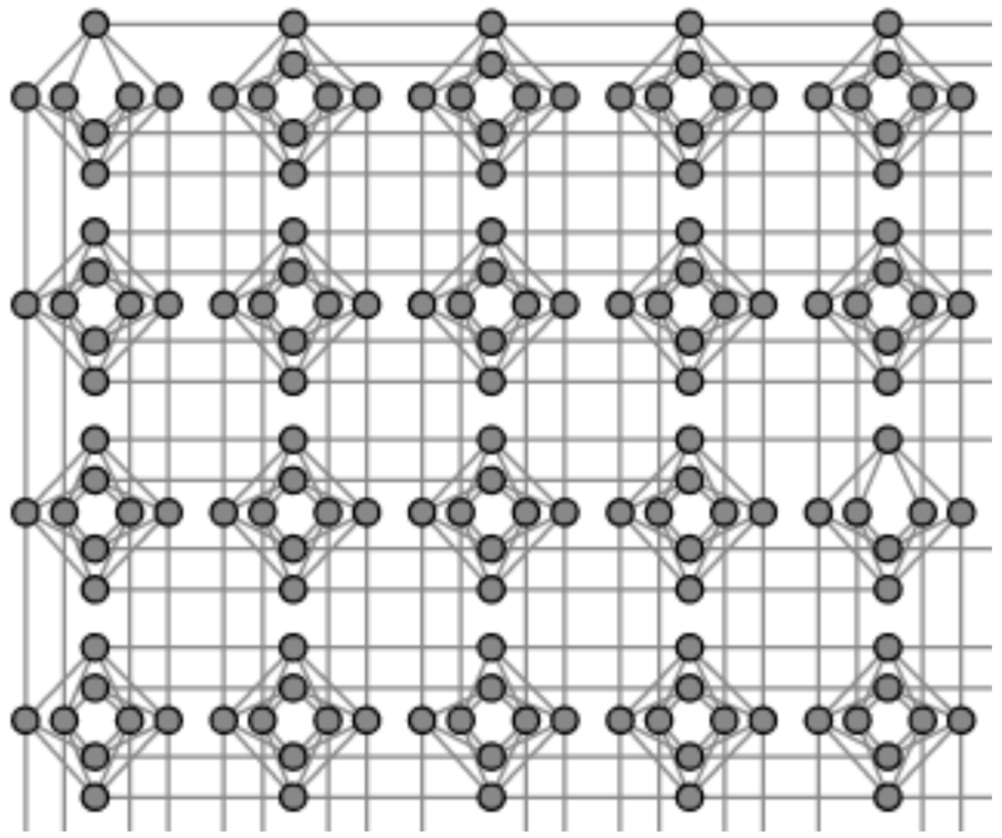

GUROBI OPTIMIZATION

Billionnet, Alain, and Sourour Elloumi. "Using a mixed integer quadratic programming solver for the unconstrained quadratic 0-1 problem." Mathematical Programming 109.1 (2007): 55-68.

HFS Local Search Solver

Hamze-de Freitas-Selby (HFS)

https://arxiv.org/pdf/1207.4149.pdf **HF**

https://github.com/alex1770/QUBO-Chimera **S**

Low Treewidth Subgraphs

Optimization Loop using DP

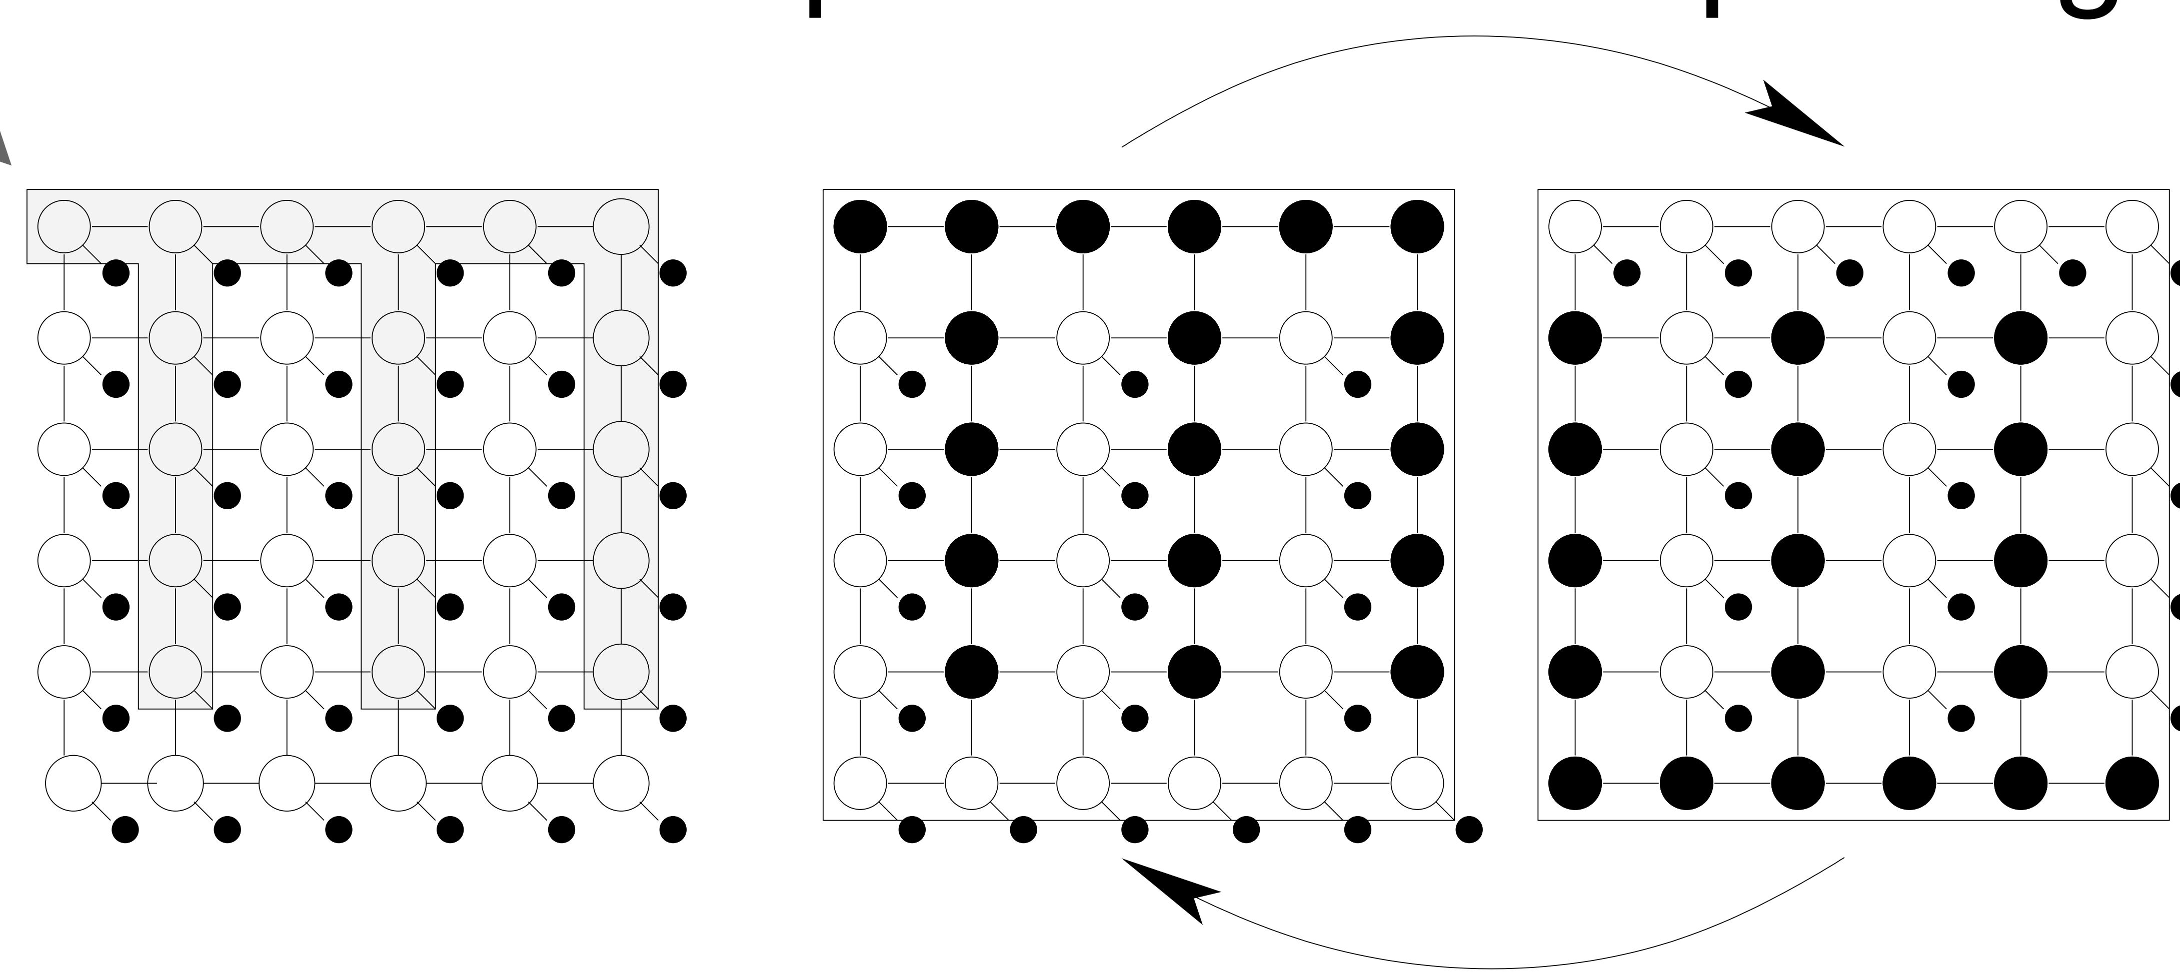

highly optimized C code

D:WAVE

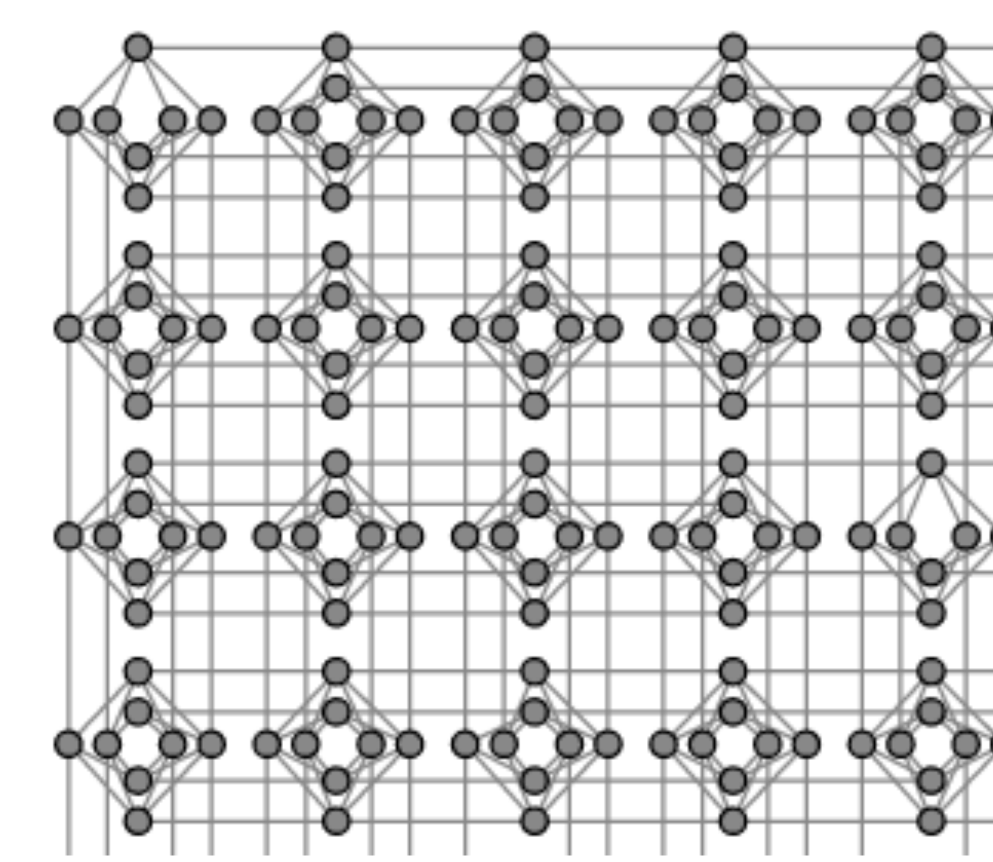

Benchmark #1: QPLIB

MI-QCQP

$$\min: \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij} x_i x_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_i x_i$$

s.t.

$$\sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{sij} x_i x_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_{si} x_i \leq 0 \ \forall s \in \mathcal{S}$$

$$x_i \in \{0,1\} \ \forall i \in \mathcal{B} \subseteq \mathcal{N}$$

$$x_i \in \mathbb{N} \ \forall i \in \mathcal{I} \subseteq \mathcal{N}$$

$$x_i \in \mathbb{R} \ \forall i \in \mathcal{R} \subseteq \mathcal{N}$$

cnrs

QPLIB

$\Rightarrow$

QUBO Model

$$\min: \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij} x_i x_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_i x_i$$

s.t.

$$x_i \in \{0,1\} \ \forall i \in \mathcal{N}$$

$$|\mathcal{N}| \leq 1100$$

$$|\mathcal{E}| \leq 3068$$

D:WAVE

Benchmark #1: QPLIB

- **479 cases**
  - -350 cases, no real or int
  - -87 cases, no constraints
  - -23 cases, edge bound
  - -2 cases, variable bound
  - **17 cases** remain…

QPLIB

Benchmark #1: QPLIB

			gurobi			dwave				
Case	$\|\mathcal{N}\|$	$\|\mathcal{E}\|$	Best Sol.	Opt. Gap	Time	Best Sol.	Best	Inf.	Samples	Time
qplib_3876	28	48	-24	0%	<1	-24	8574	352	10000	0+4
qplib_3607	66	120	-68	0%	<1	-68	1720	4424	10000	1+4
qplib_5727	225	450	-15051133	0%	<1	F.E.	-	-	-	T.L.
qplib_3756	153	288	-160	0%	1.2	-160	234	8451	10000	57+5
qplib_5755	400	800	-24838942	0%	1.5	F.E.	-	-	-	T.L.
qplib_3821	190	360	-192	0%	15	F.E.	-	-	-	T.L.
qplib_3565	276	528	-282	0%	93	F.E.	-	-	-	T.L.
qplib_3705	378	728	-384	0%	202	F.E.	-	-	-	T.L.
qplib_3745	325	624	-334	0%	753	F.E.	-	-	-	T.L.
qplib_3506	496	960	-478	1%	T.L.	F.E.	-	-	-	T.L.
qplib_3738	435	840	-422	1%	T.L.	F.E.	-	-	-	T.L.
qplib_3877	630	1224	-602	1%	T.L.	F.E.	-	-	-	T.L.
qplib_3706	703	1368	-682	2%	T.L.	F.E.	-	-	-	T.L.
qplib_3642	1035	2024	-1030	3%	T.L.	F.E.	-	-	-	T.L.
qplib_3650	946	1848	-918	3%	T.L.	F.E.	-	-	-	T.L.
qplib_5889	250	3045	-40358	15%	T.L.	F.E.*	-	-	-	T.L.
qplib_5909	250	3015	-33587	28%	T.L.	F.E.*	-	-	-	T.L.

F.E.

Failed Embed

T.L.

Time Limit

(1 hour)

Benchmark #2: DIMACS Max-Clique

http://iridia.ulb.ac.be/~fmascia/maximum_clique/DIMACS-benchmark

dimacs benchmark set

This is a selected set of instances form the Second DIMACS Implementation Challenge (1992-1993).
In the following table, the clique number ω(G) corresponds to the global optimum or to the lower bound as indicated by the
For some instances the lower bound has been confirmed to coincide with the global optimum. For such instances, the cliqu
The exact algorithms that confirmed the bound are reported later on this page.

instance	ω(g)	best known	nodes	edges	graph degrees		best degrees	
					median	iqr	median	iqr
C125.9	34*	34	125	6 963	112.0	(5.00)	114.5	(4.75)
C250.9	44*	44	250	27 984	224.0	(6.00)	227.0	(5.00)
C500.9	≥ 57	57	500	112 332	449.0	(9.00)	455.0	(9.00)
C1000.9	≥ 68	68	1 000	450 079	900.0	(13.00)	907.0	(11.25)
C2000.9	≥ 80	80	2 000	1 799 532	1 800.0	(18.00)	1 803.0	(15.25)
DSJC1000_5	15	15	1 000	499 652	500.0	(20.00)	503.0	(23.00)
DSJC500_5	13	13	500	125 248	250.0	(16.00)	259.0	(14.00)
C2000.5	16*	16	2 000	999 836	999.0	(30.00)	1 006.0	(11.50)

Why Max-Clique?

- **Formulate on the Compliment Graph**
  - super sparse, i.e. easy to embed!
- **Combinatorial Problem**
  - minimal issues due to coefficient accuracy
- **Easy to generate problems**

Benchmark #2: DIMACS Max-Clique

			gurobi			dwave				
Case	$\lvert\mathcal{N}\rvert$	$\lvert\mathcal{E}\rvert$	Best Sol.	Opt. Gap	Time	Best Sol.	Best	Inf.	Samples	Time
C015_9	15	12	-11	0%	<1	-11	9073	1	10000	0+3
C020_9	20	17	-14	0%	<1	-14	8370	85	10000	0+3
C030_9	30	44	-16	0%	<1	-16	5651	123	10000	0+3
C040_9	40	77	-18	0%	<1	-18	3865	316	10000	0+4
C050_9	50	108	-24	0%	<1	-24	16	1254	10000	0+4
C060_9	60	158	-25	0%	<1	-25	22	5465	10000	0+5
C070_9	70	215	-27	0%	<1	-26	1	9855	10000	4+5
C080_9	80	306	-29	0%	<1	F.E.	-	-	-	T.L.
C090_9	90	407	-29	0%	1.0	F.E.	-	-	-	T.L.
C100_9	100	508	-30	0%	2.0	F.E.	-	-	-	T.L.
C110_9	110	615	-32	0%	5.1	F.E.	-	-	-	T.L.
C120_9	120	729	-32	0%	45	F.E.	-	-	-	T.L.
C125_9	125	787	-34	0%	55	F.E.	-	-	-	T.L.
C250_9	250	3141	-43	40%	T.L.	F.E.*	-	-	-	T.L.

F.E.

Failed Embed

T.L.

Time Limit (1 hour)

Back to the Drawing Board

Benchmarking Stopgap

- Embedding real problems is just too difficult, at the moment.

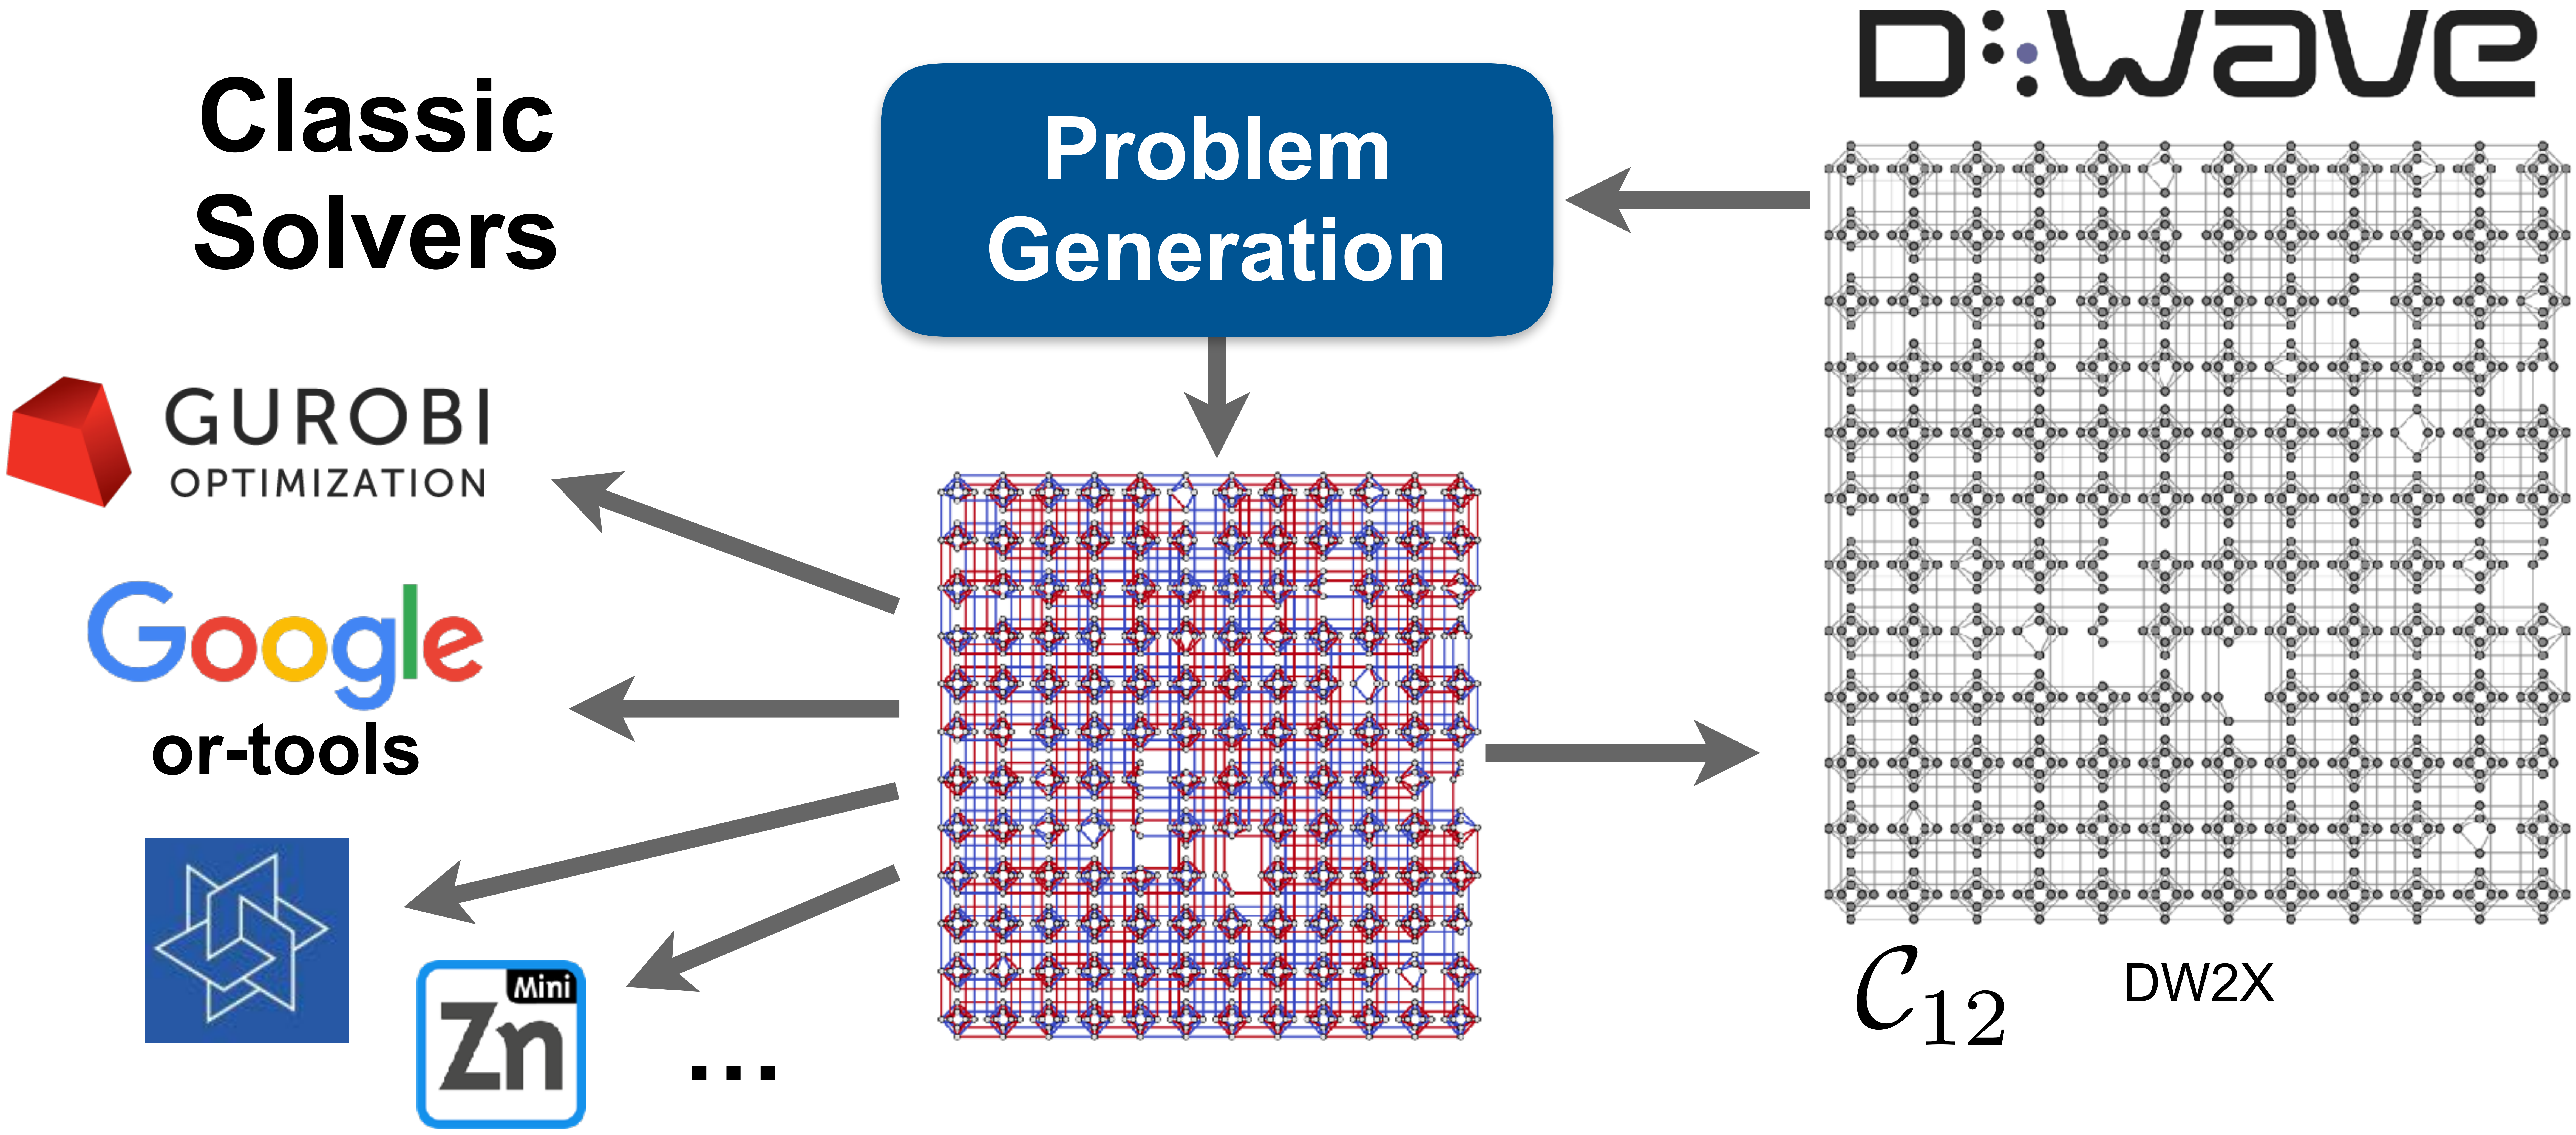

Benchmarking Tools

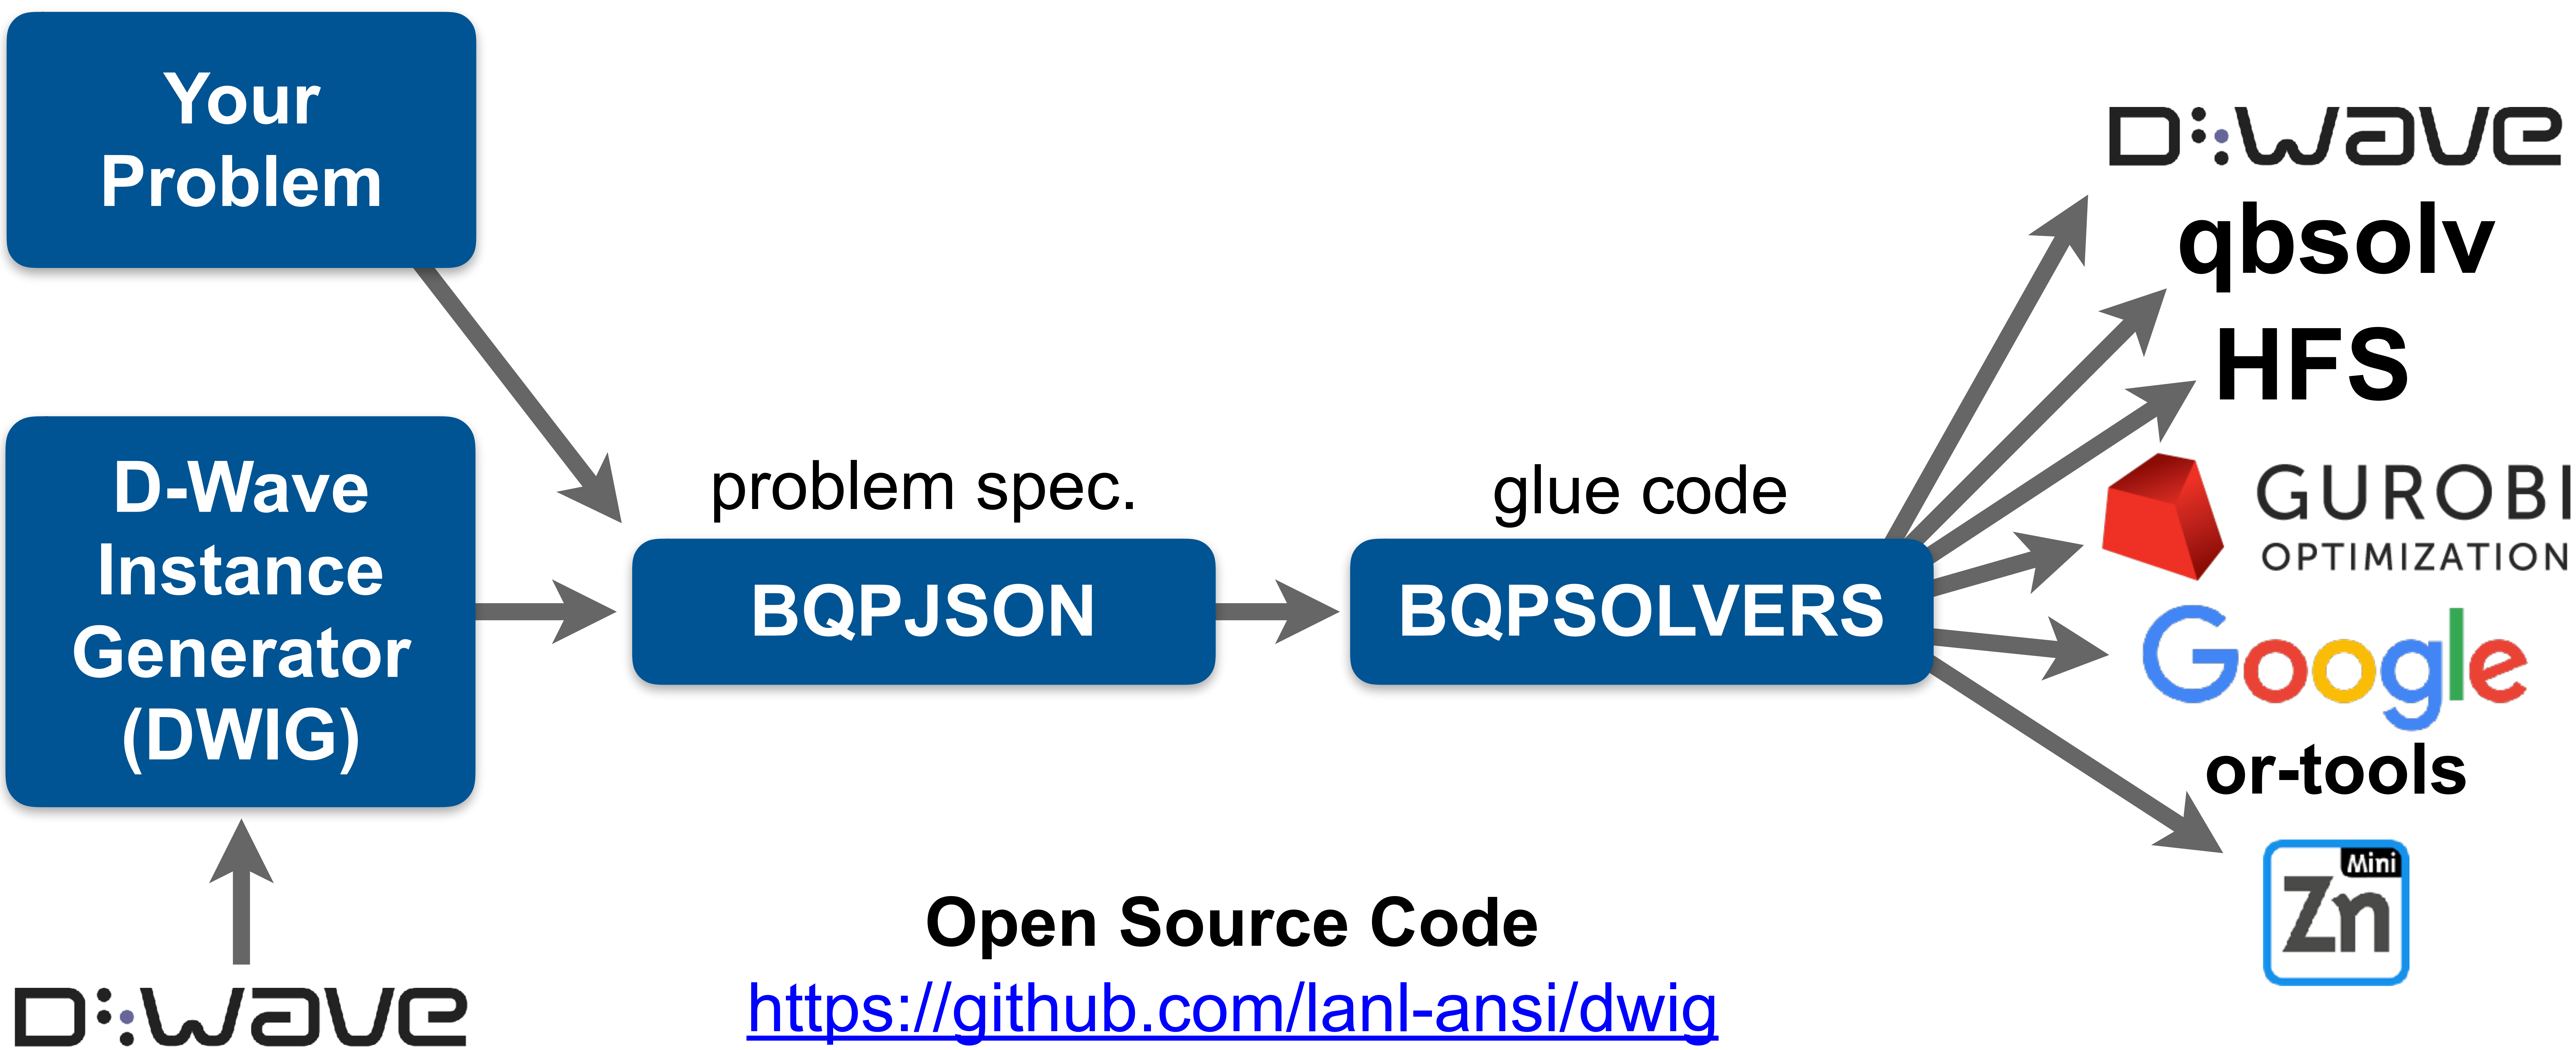

Open Source Code

https://github.com/lanl-ansi/dwig

https://github.com/lanl-ansi/bqpjson

https://github.com/lanl-ansi/bqpsolvers

The Problem with Problem Generation

- Generating interesting problems can be very hard!

Hard and Easy Distributions of SAT Problems

David Mitchell
Dept. of Computing Science
Simon Fraser University

Bart Selman
AT&T Bell Laboratories
Murray Hill, NJ 07974

Hector Levesque*
Dept. of Computer Science
University of Toronto

AAAI-92

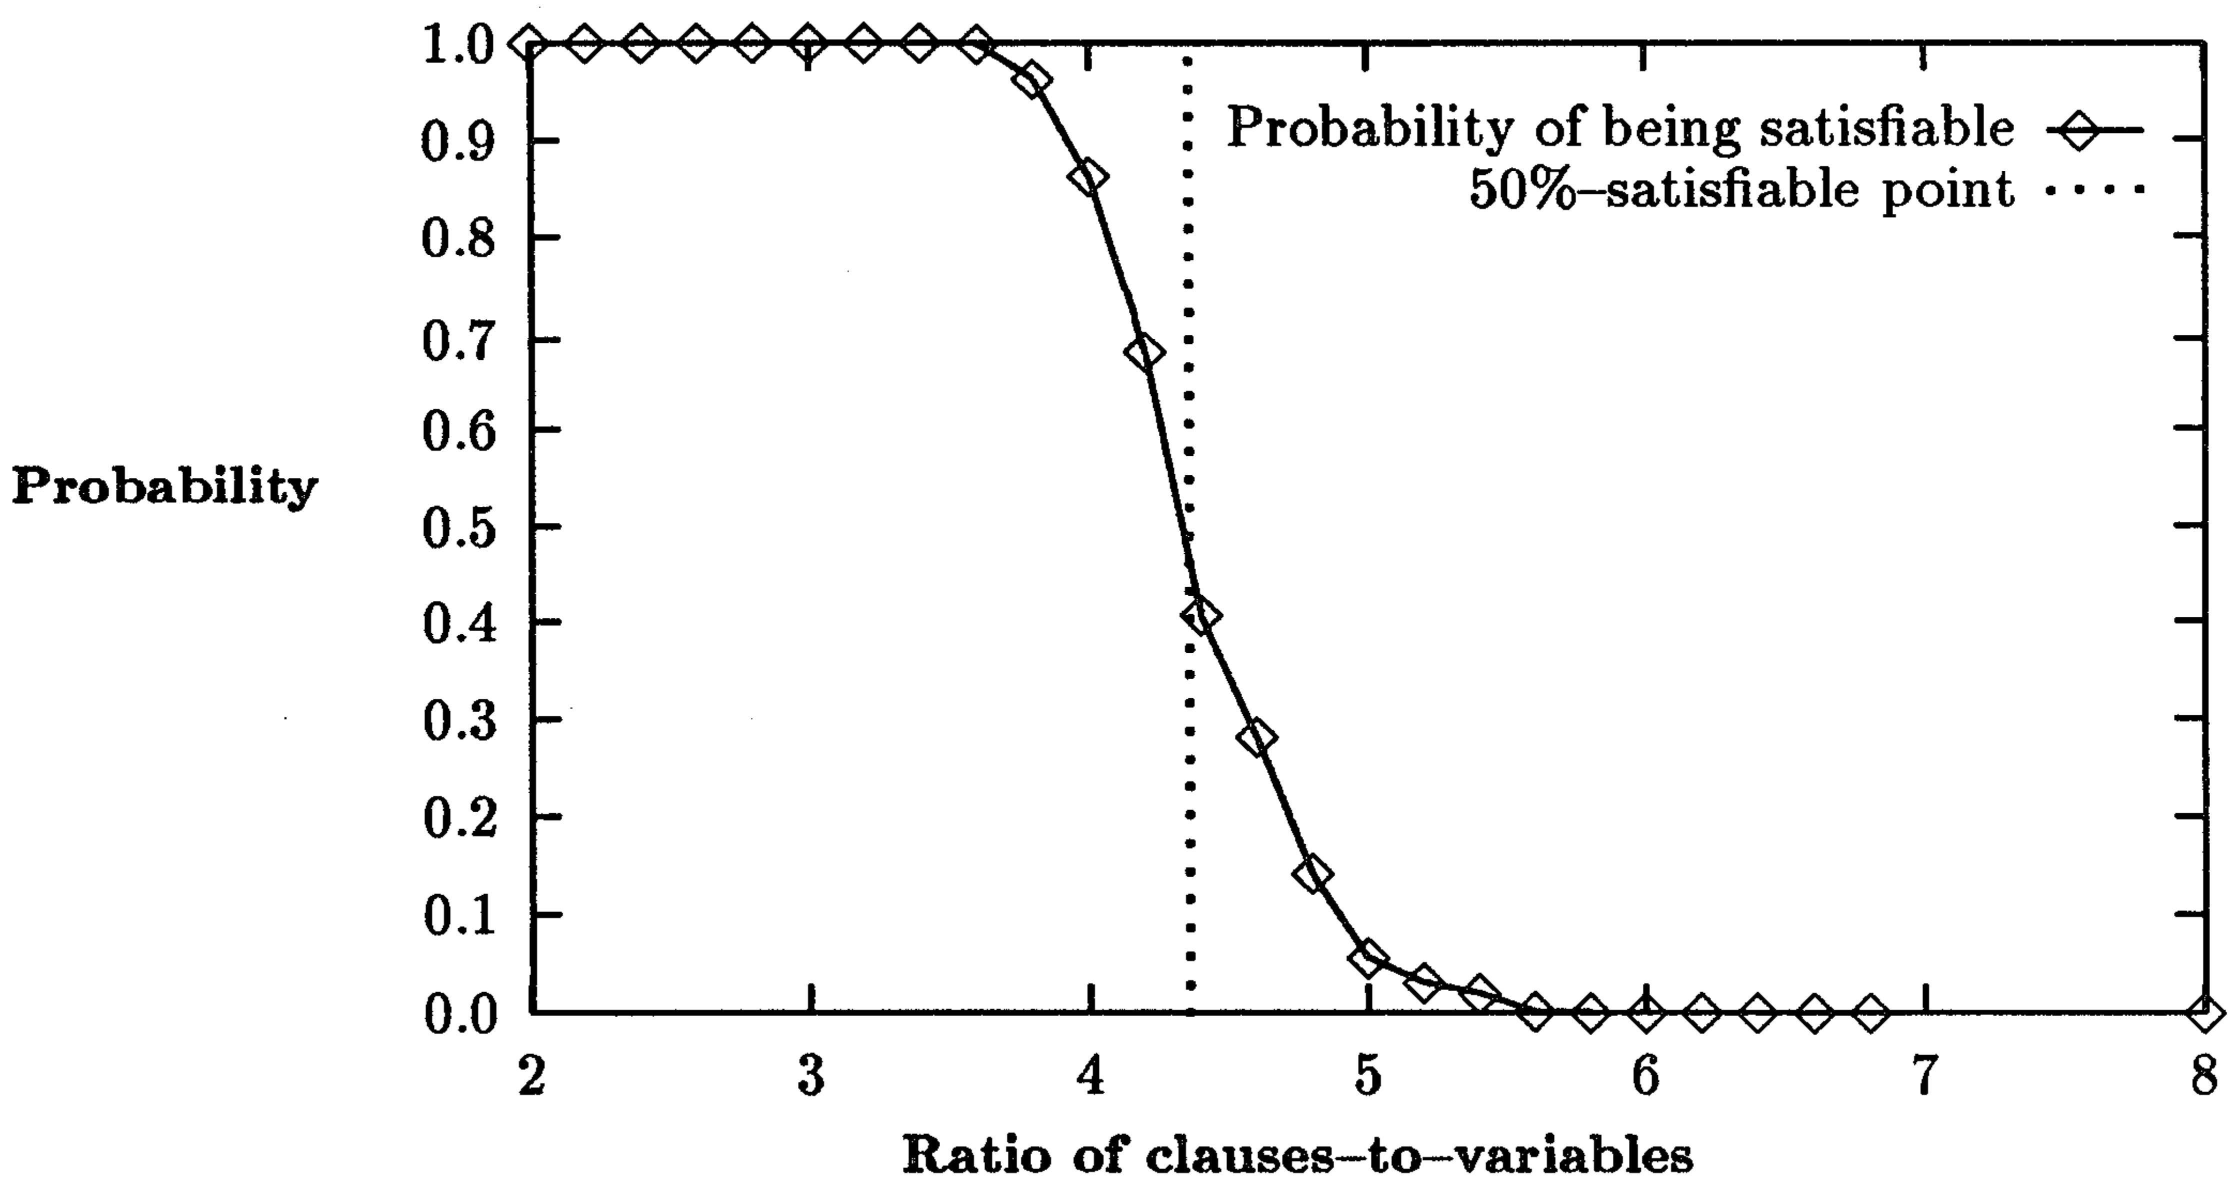

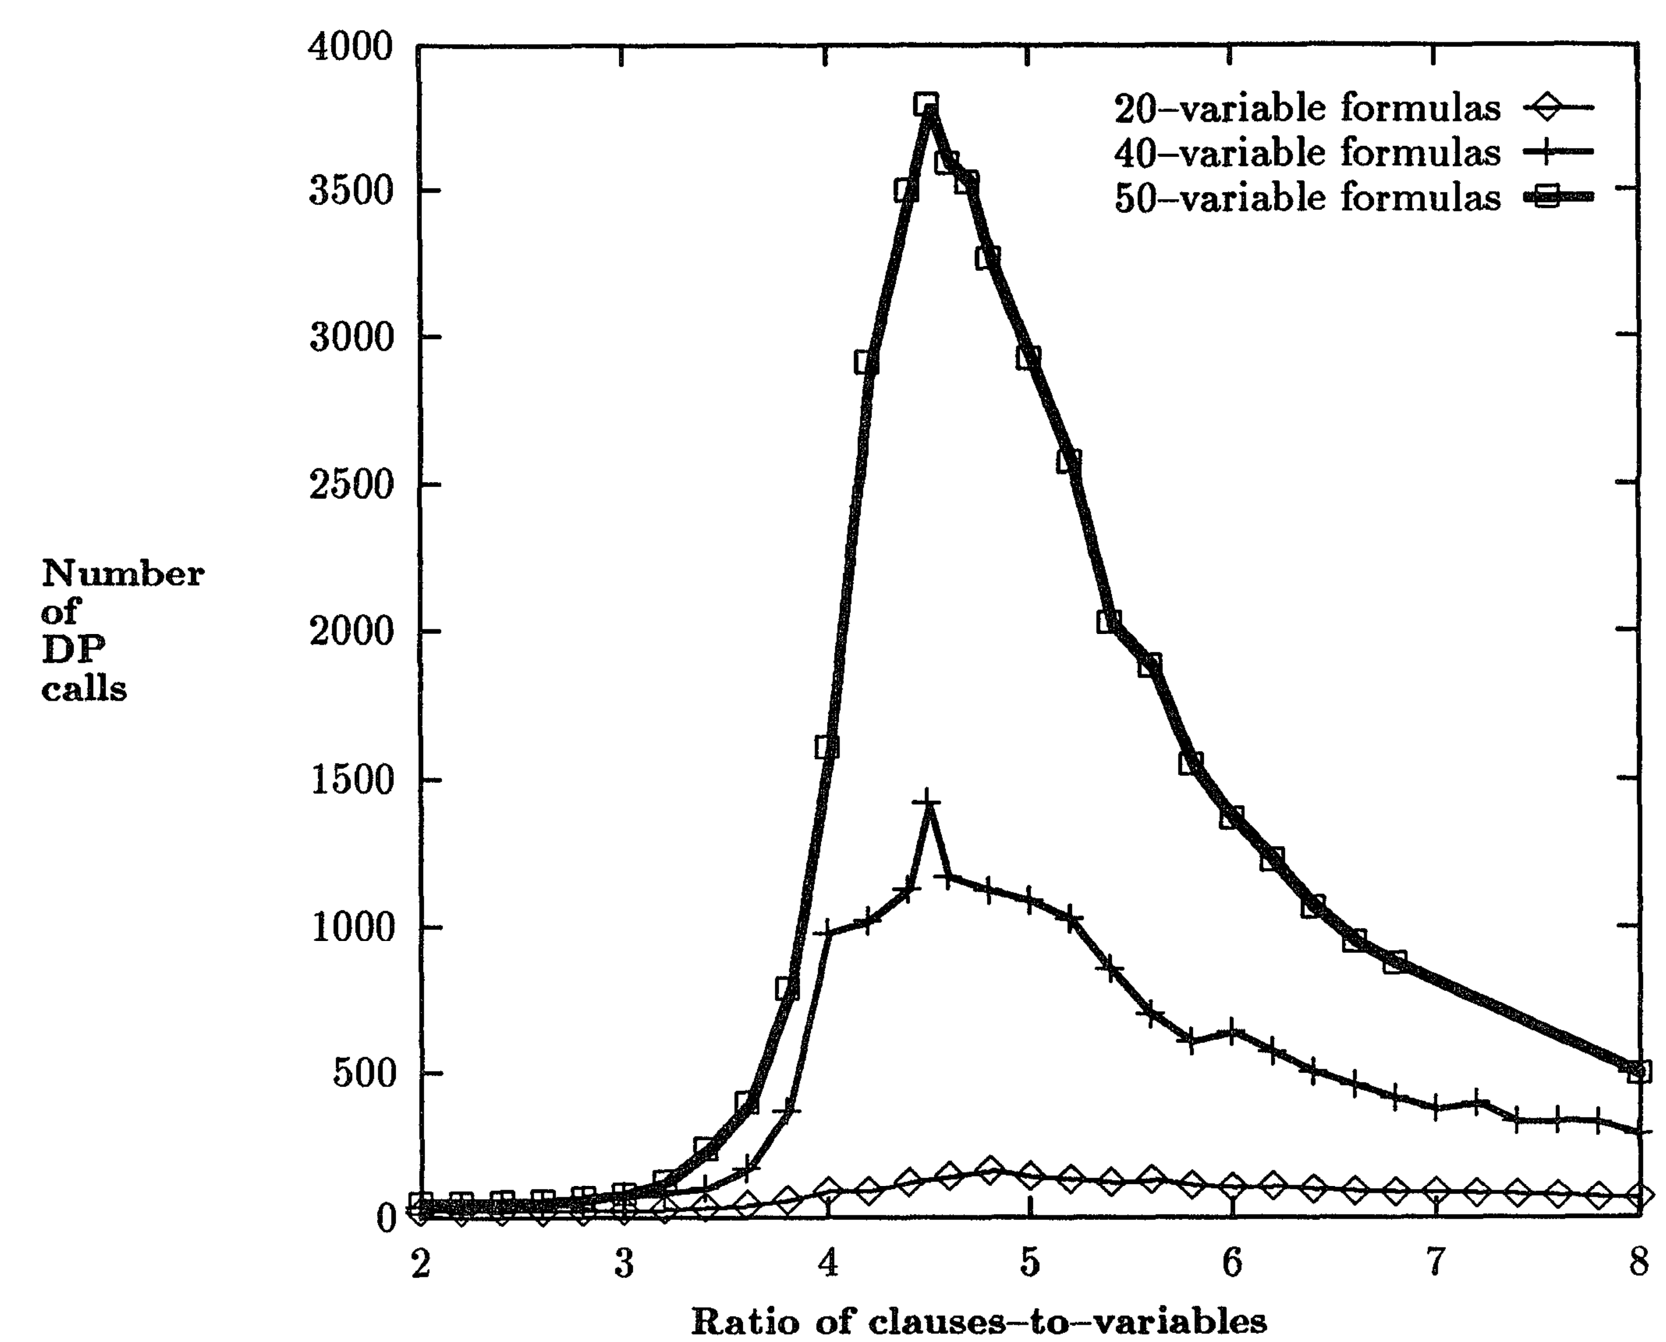

What Problems to Consider

- **Try everything from the literature**

- **RAN-k, RANF-k**
  - https://arxiv.org/abs/1508.05087

- **Frustrated Loops**
  - https://arxiv.org/abs/1508.05087
  - https://arxiv.org/abs/1701.04579
  - https://arxiv.org/abs/1703.00622

- **Weak-Strong Cluster Networks**
  - https://arxiv.org/abs/1512.02206

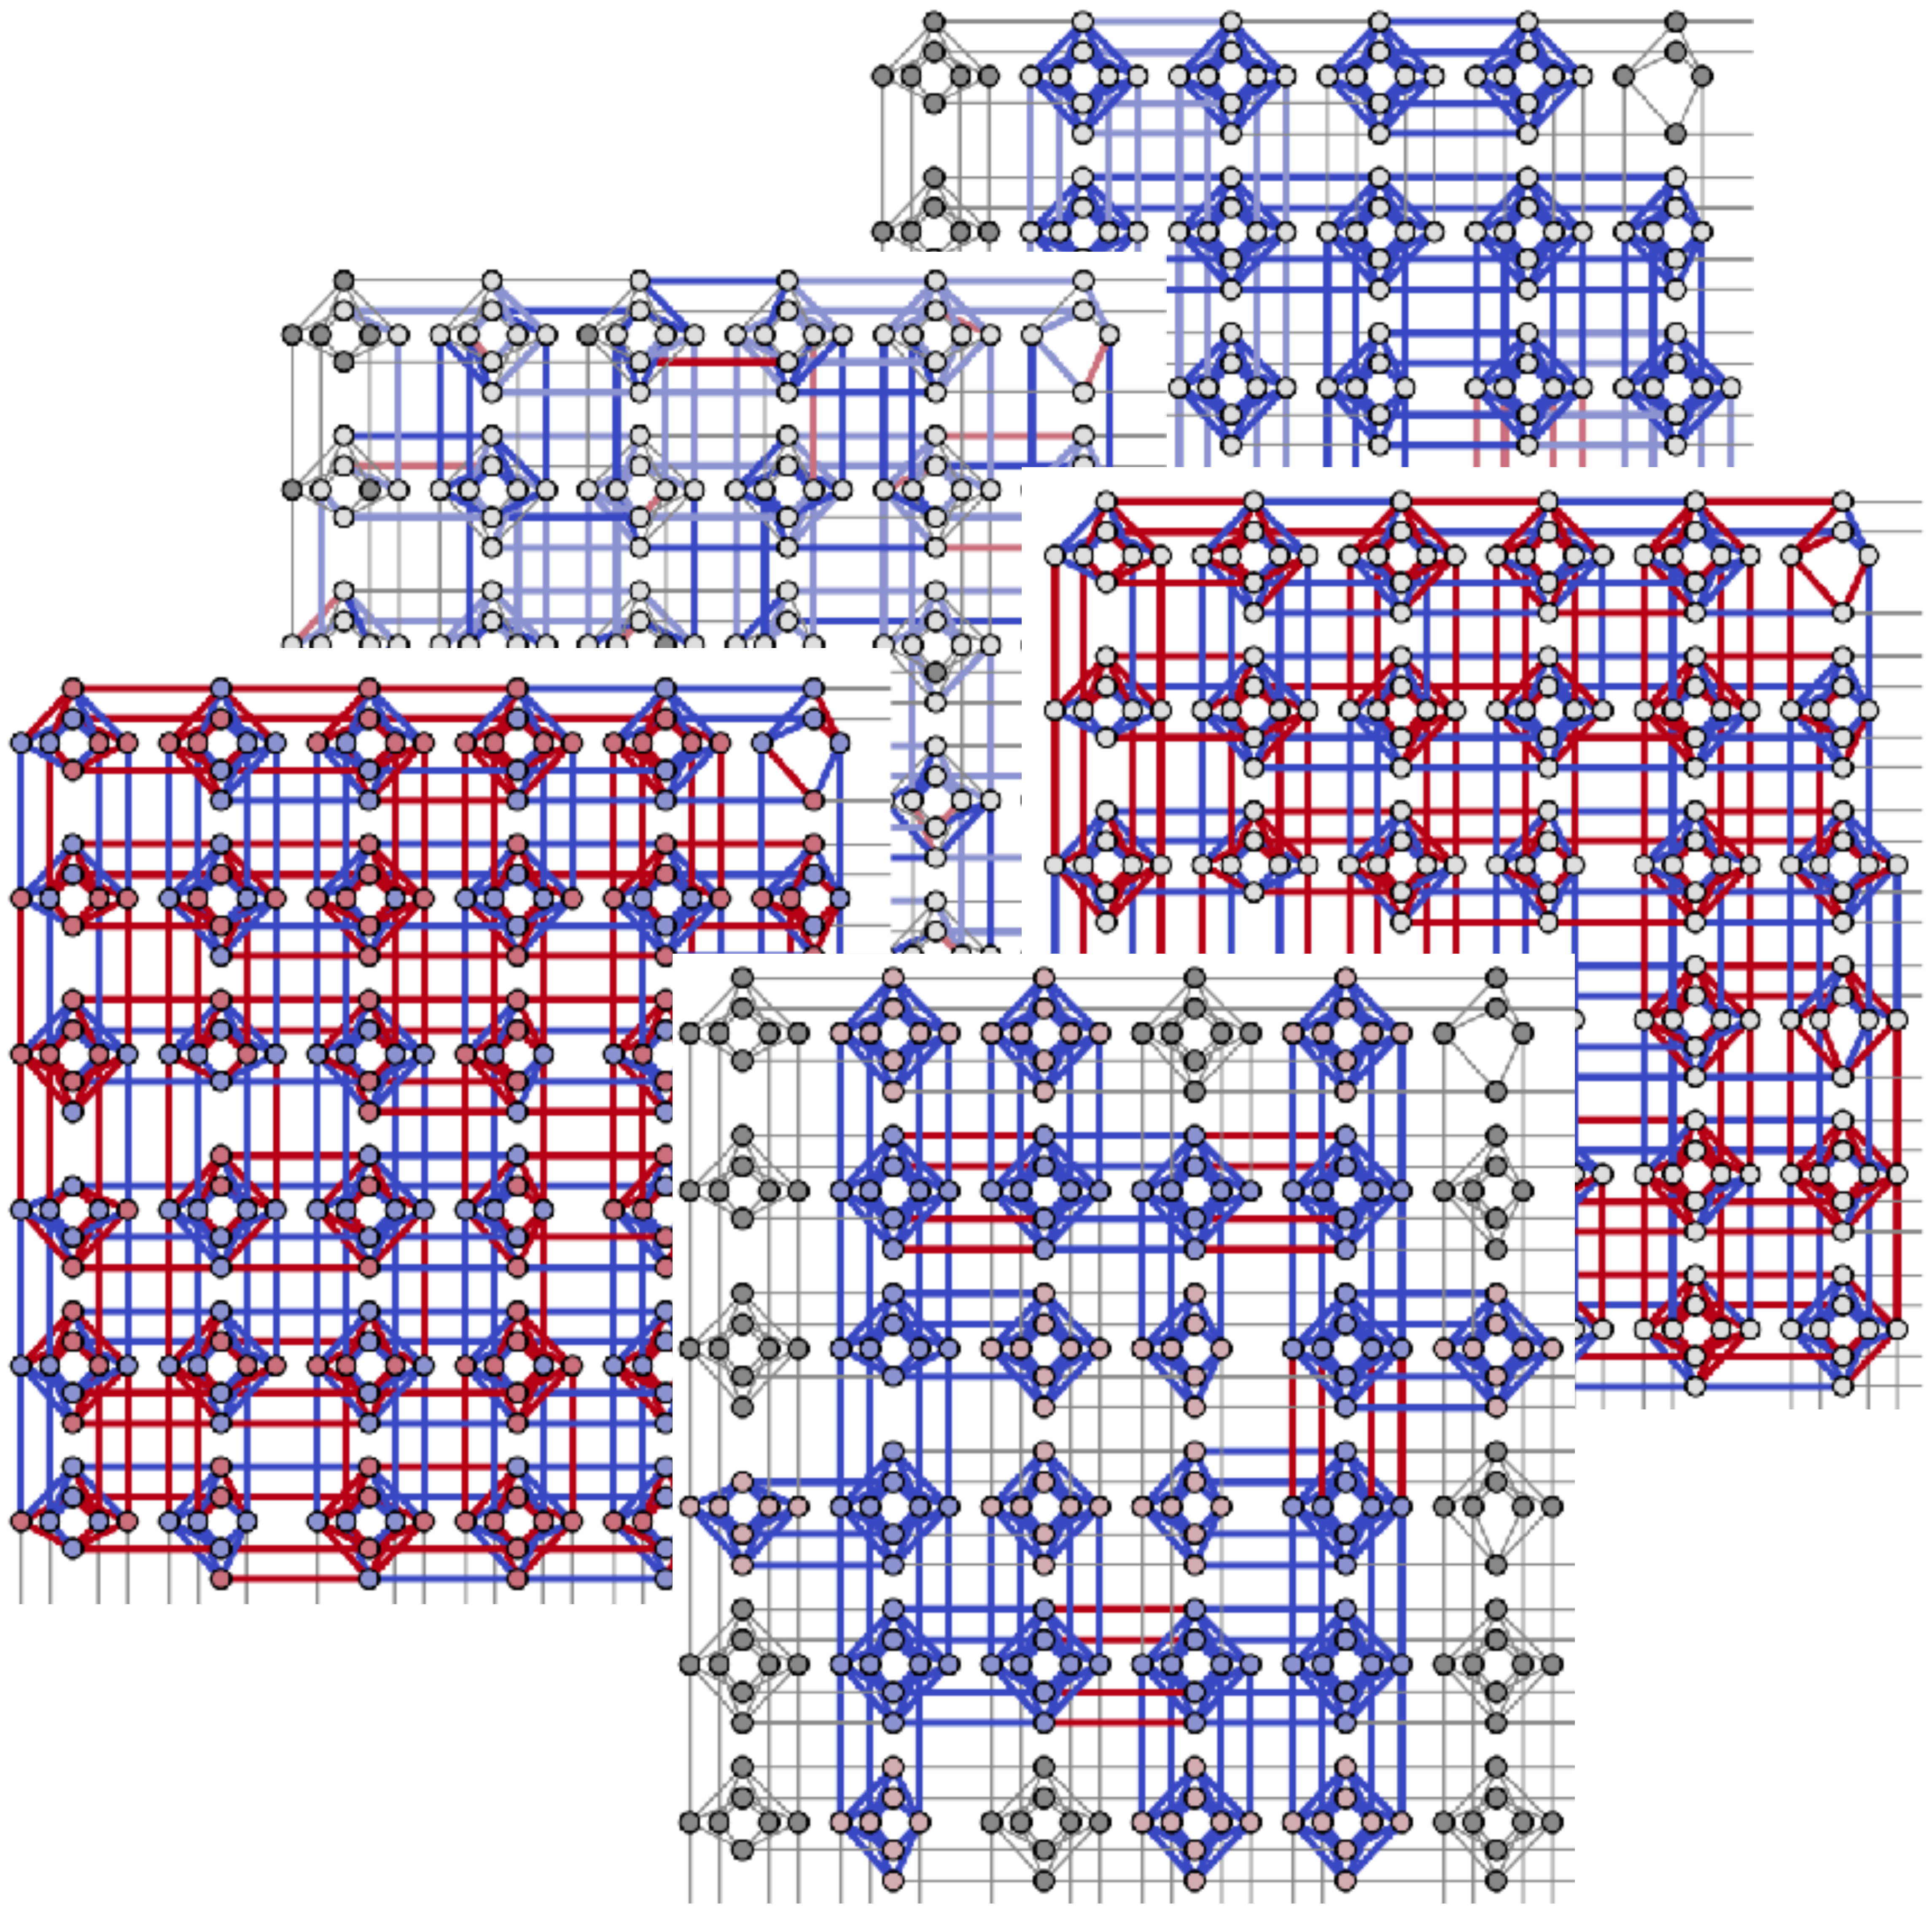

Too Easy Filtering

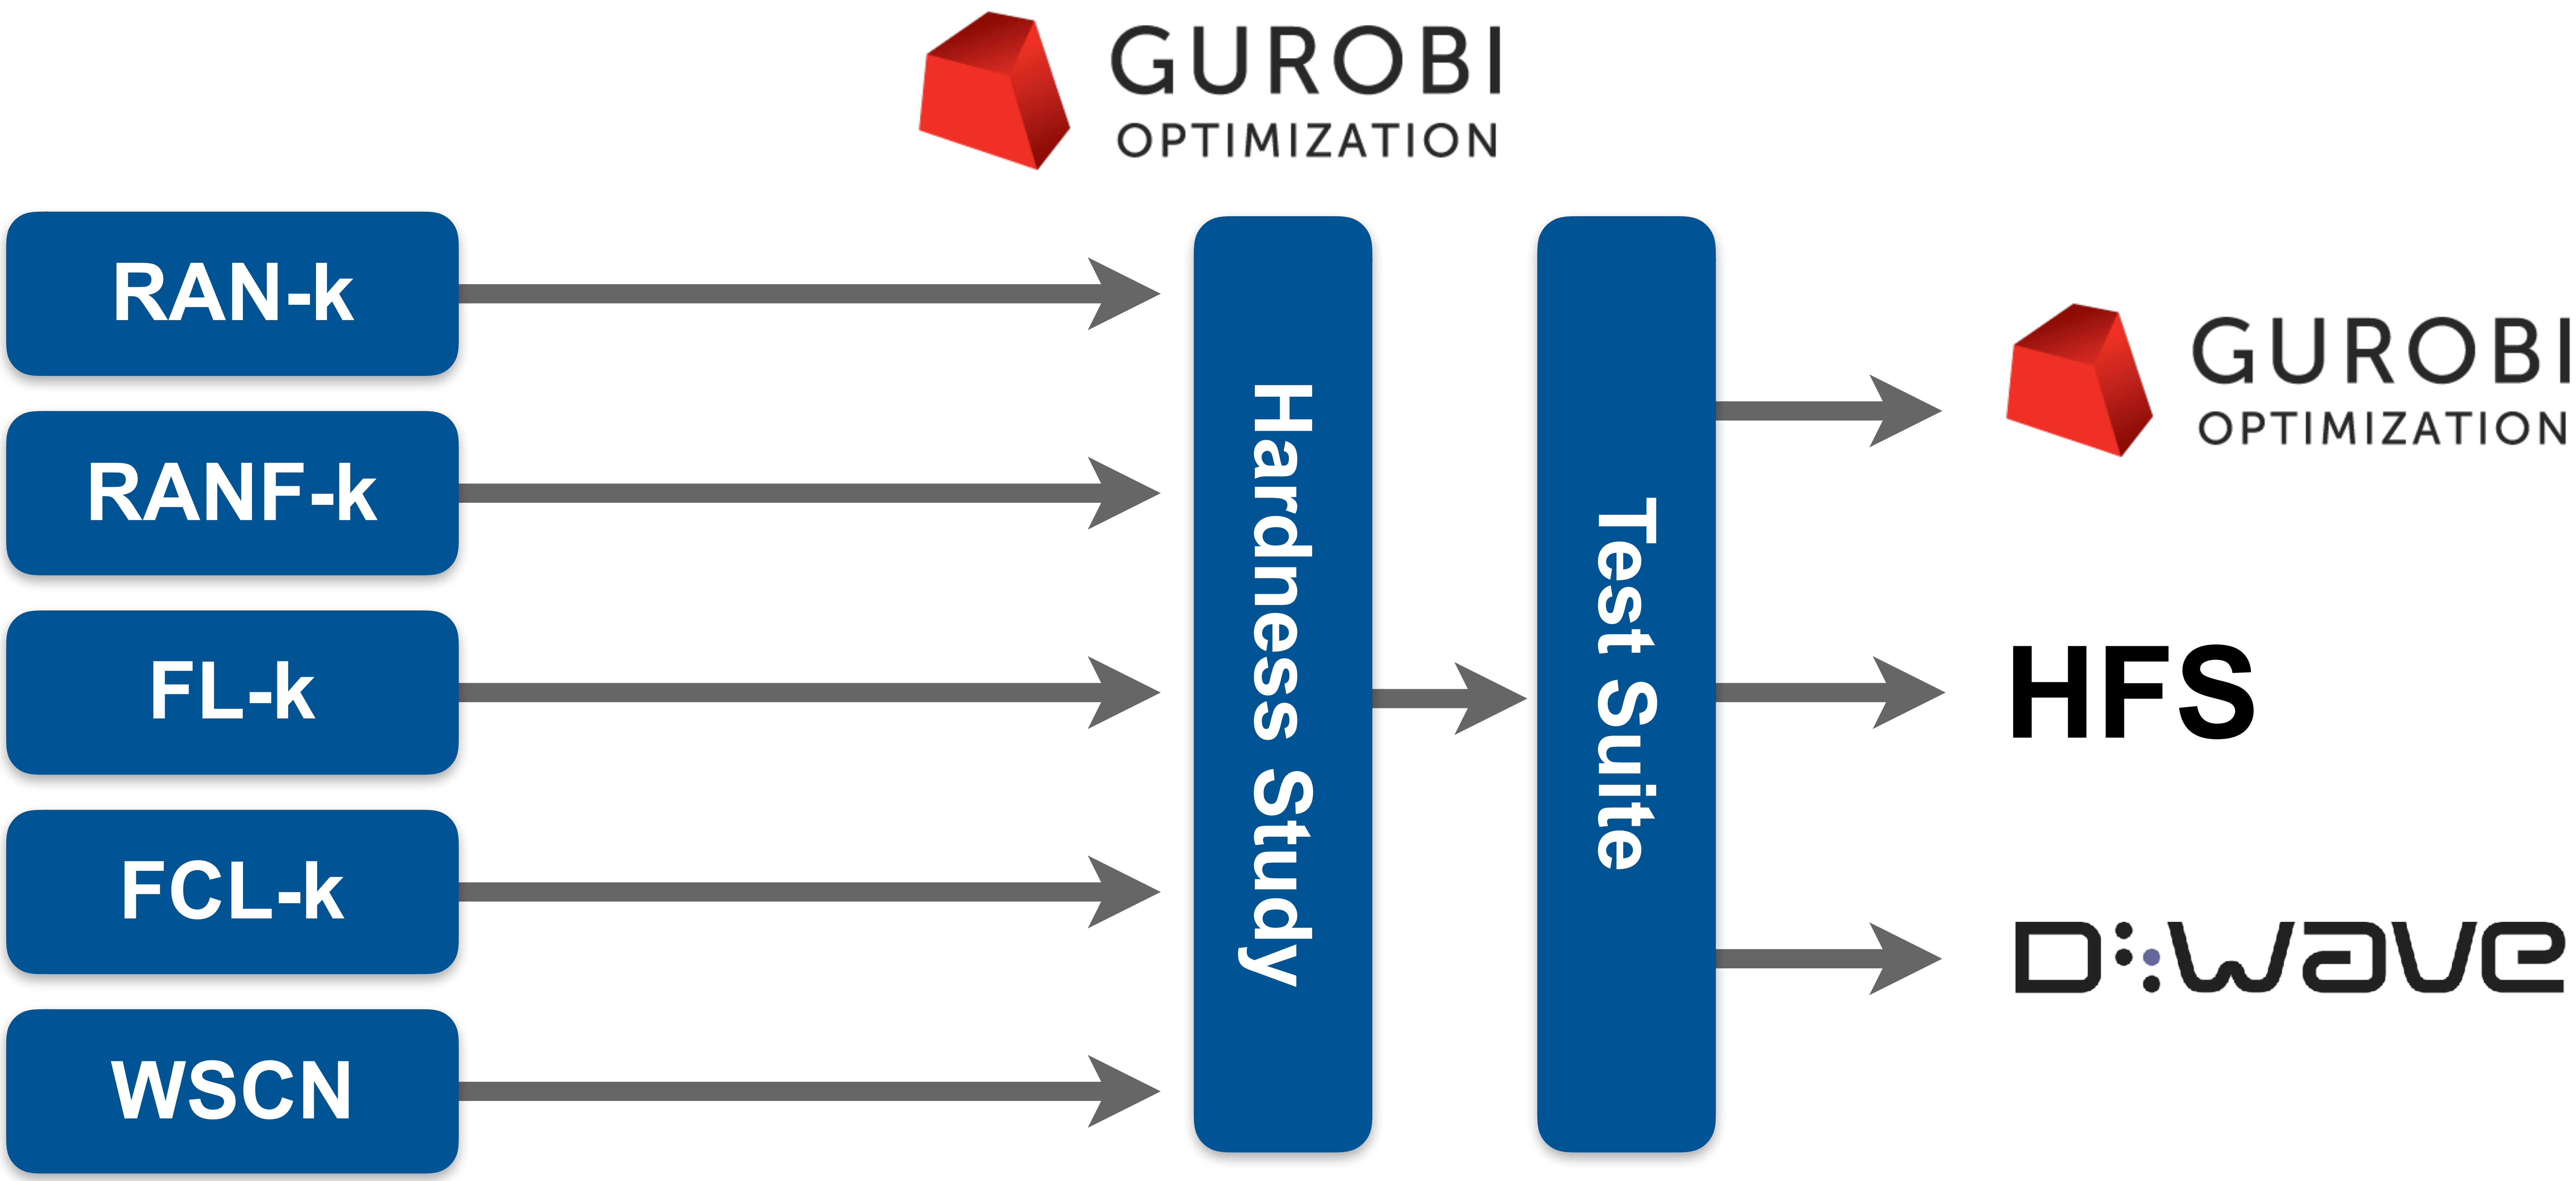

Problem Hardness

Runtime Distribution Color Coded by Problem

40,000+ Cases

What about the values of k?

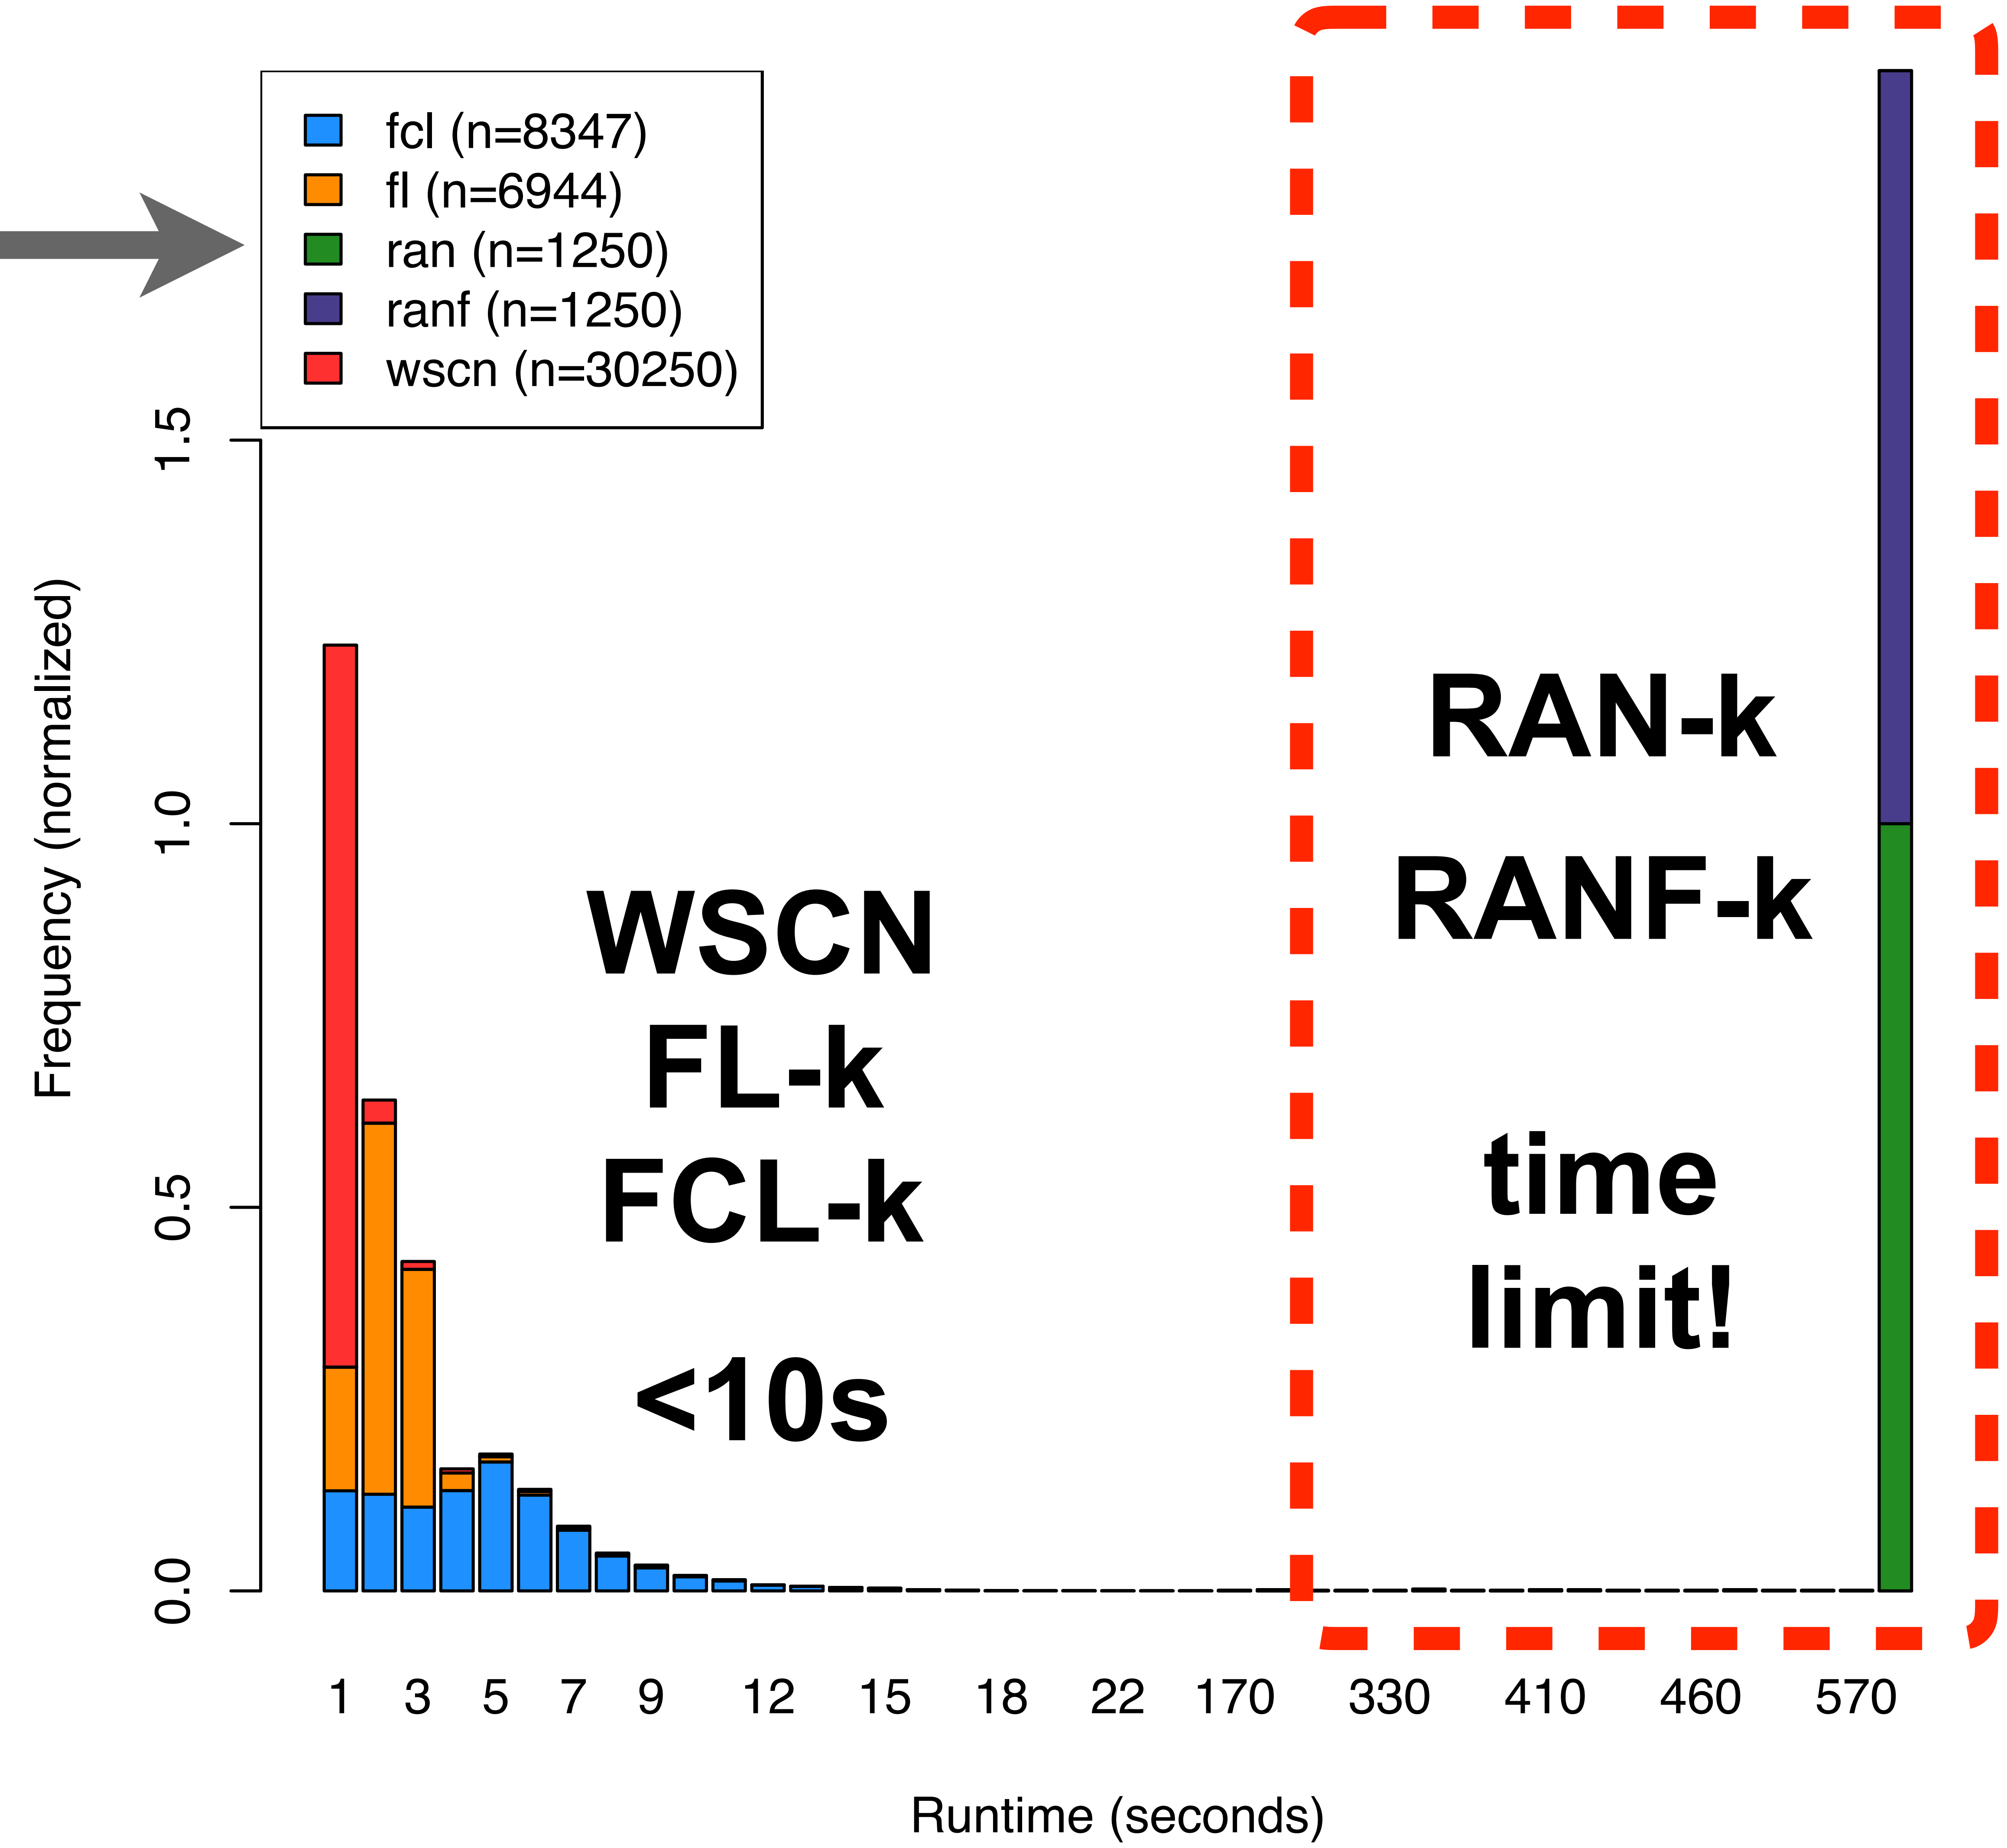

GUROBI OPTIMIZATION

600 second time limit

$\mathcal{C}_{12}$ Max Size

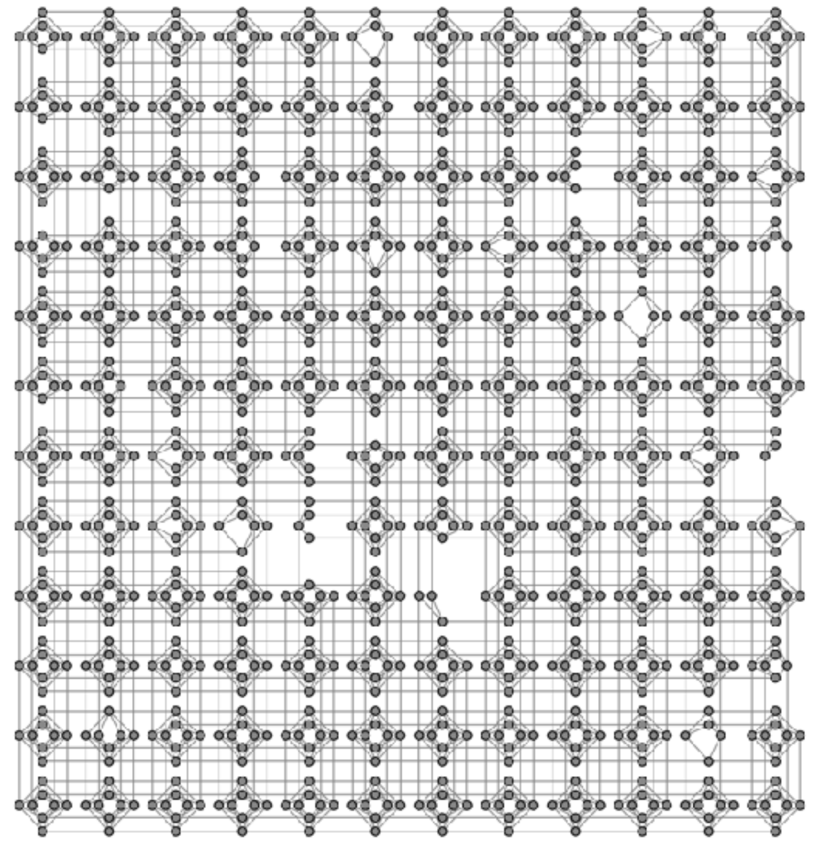

Problem Hardness

RAN-k

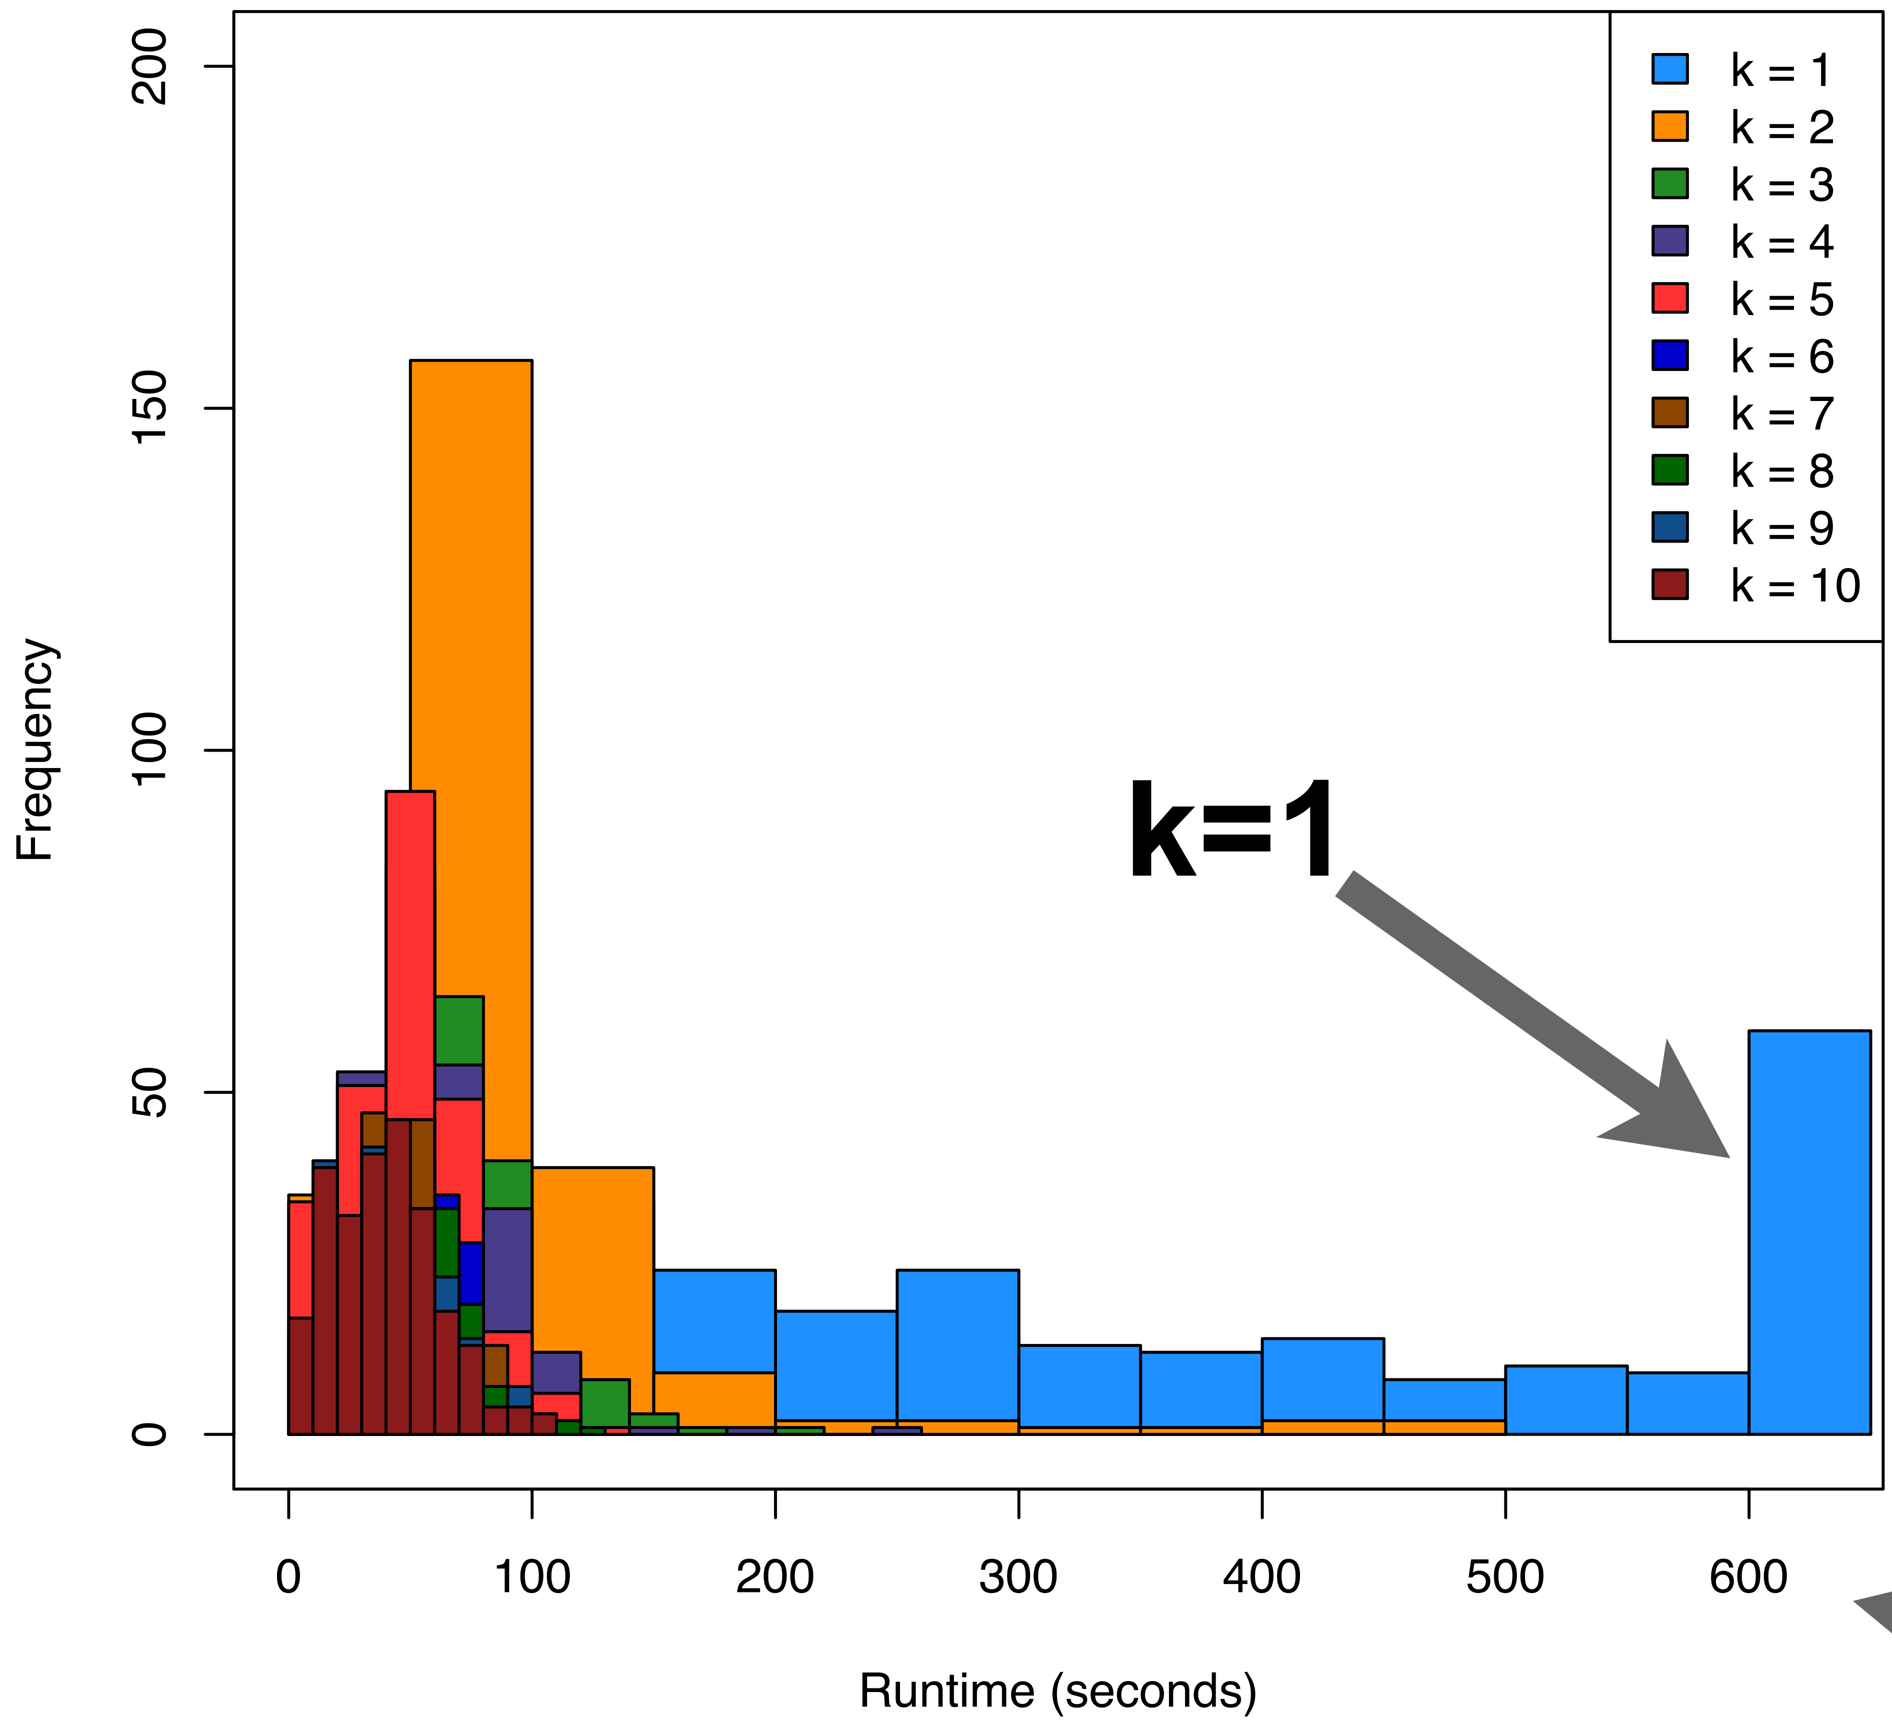

RANF-k

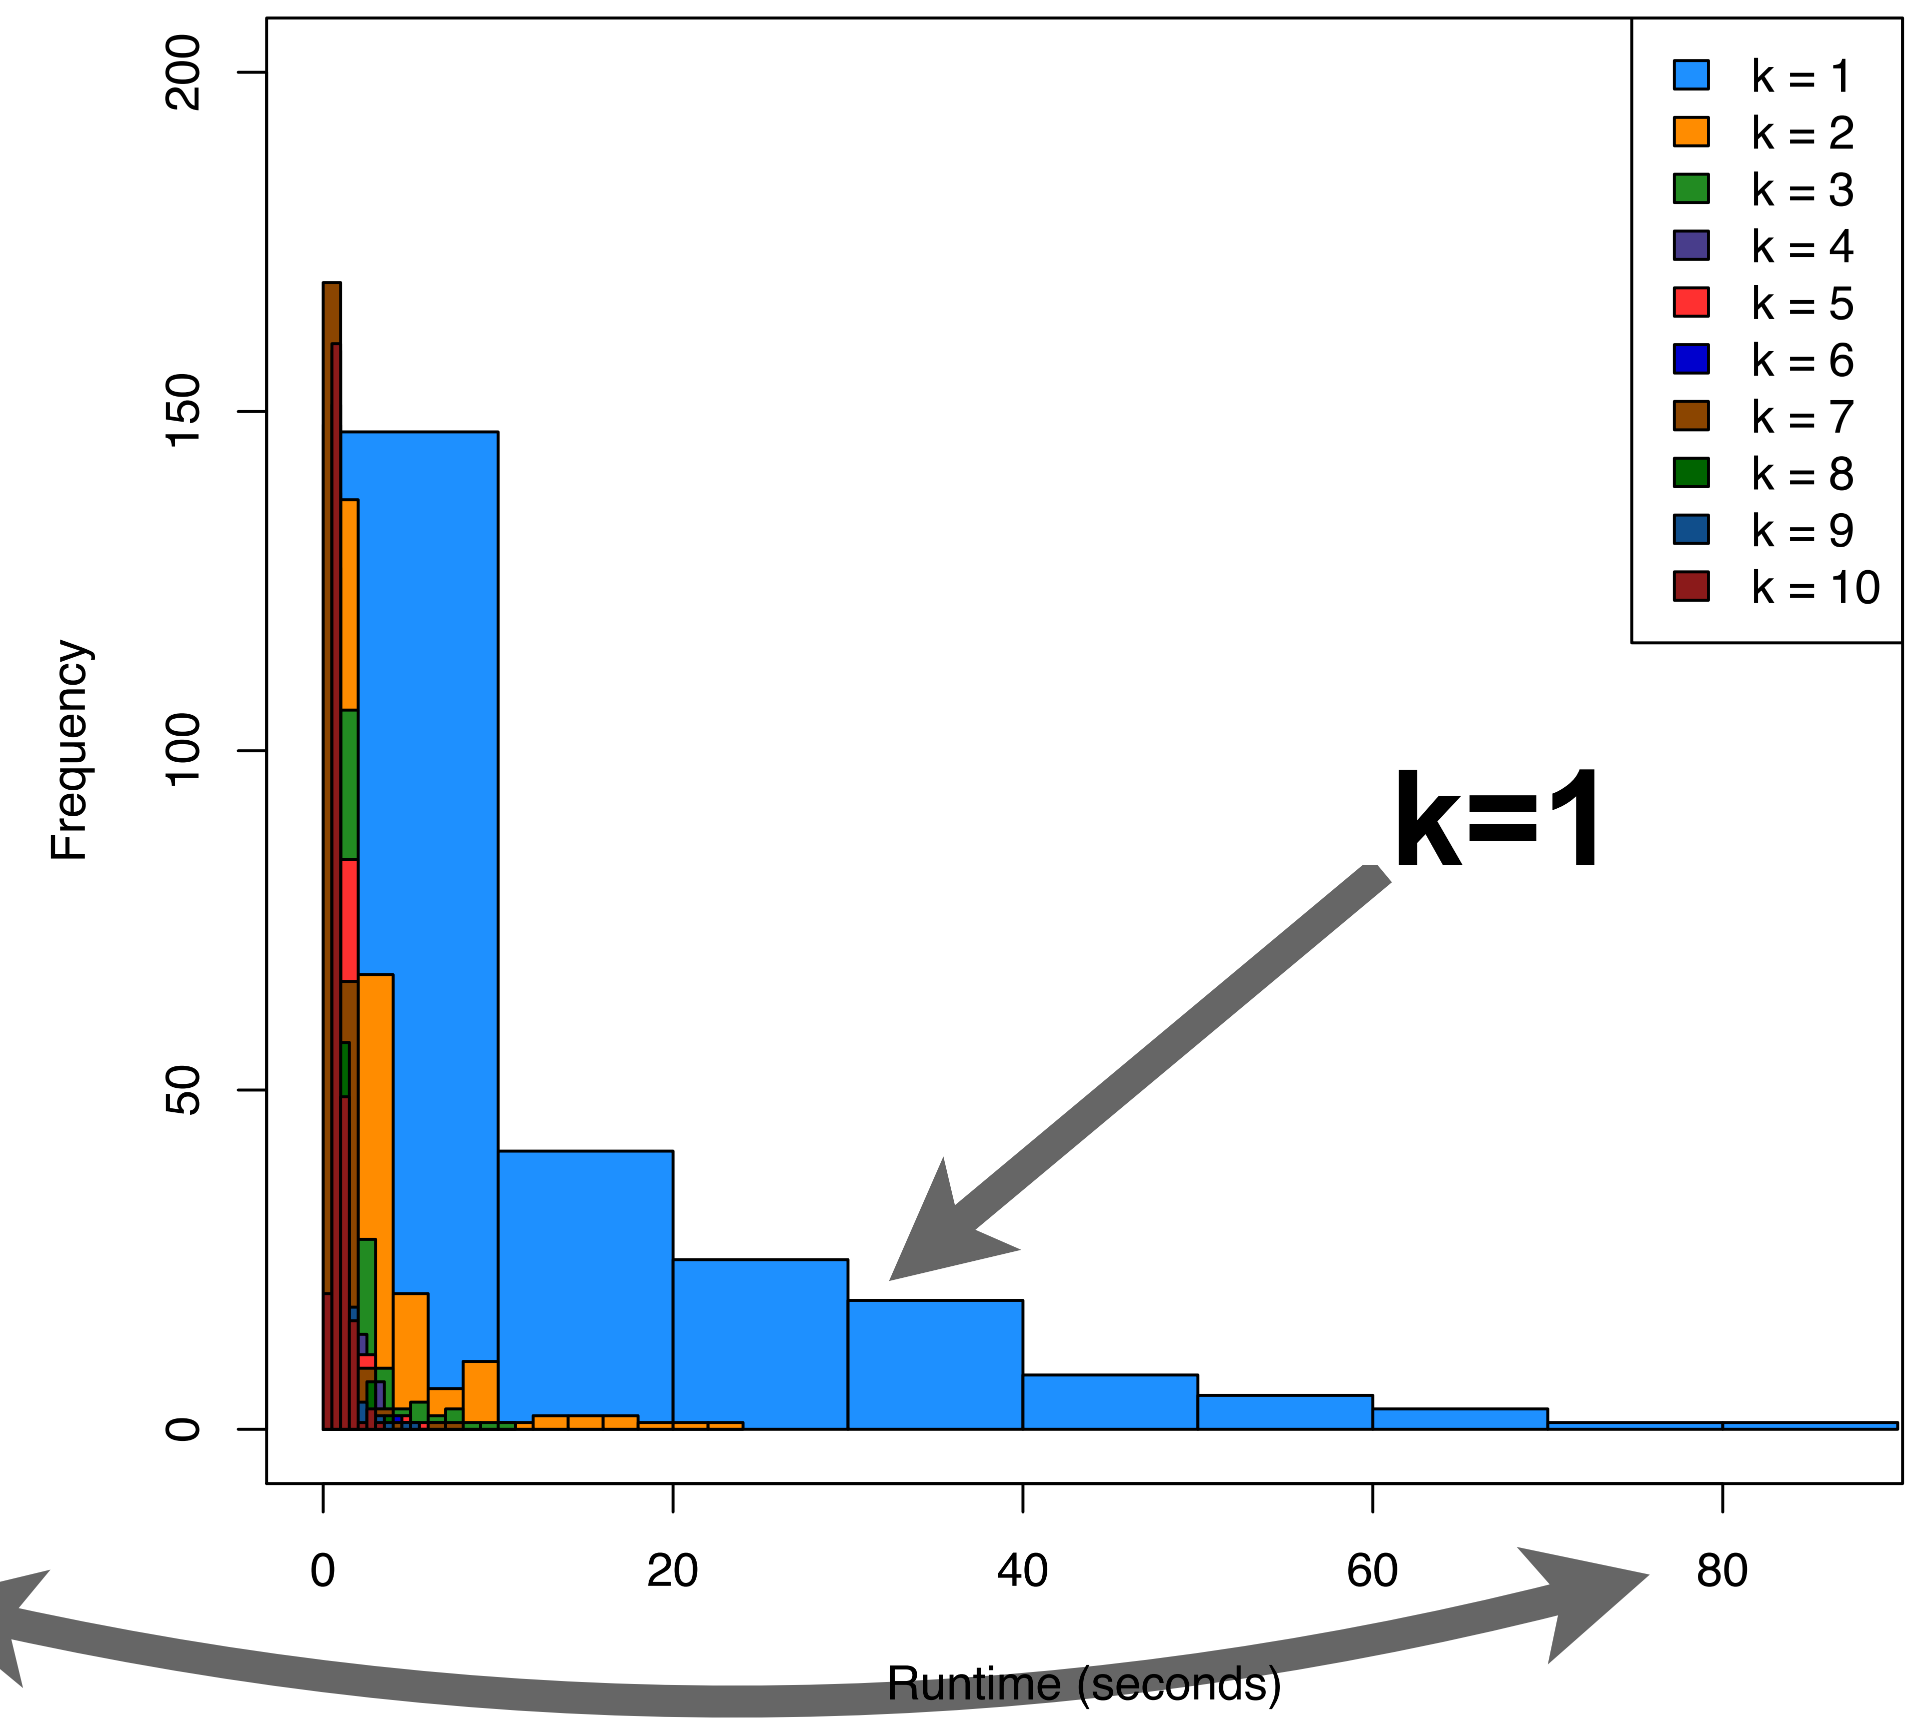

RAN > RANF

GUROBI OPTIMIZATION

optimality proof in 600 sec. or less

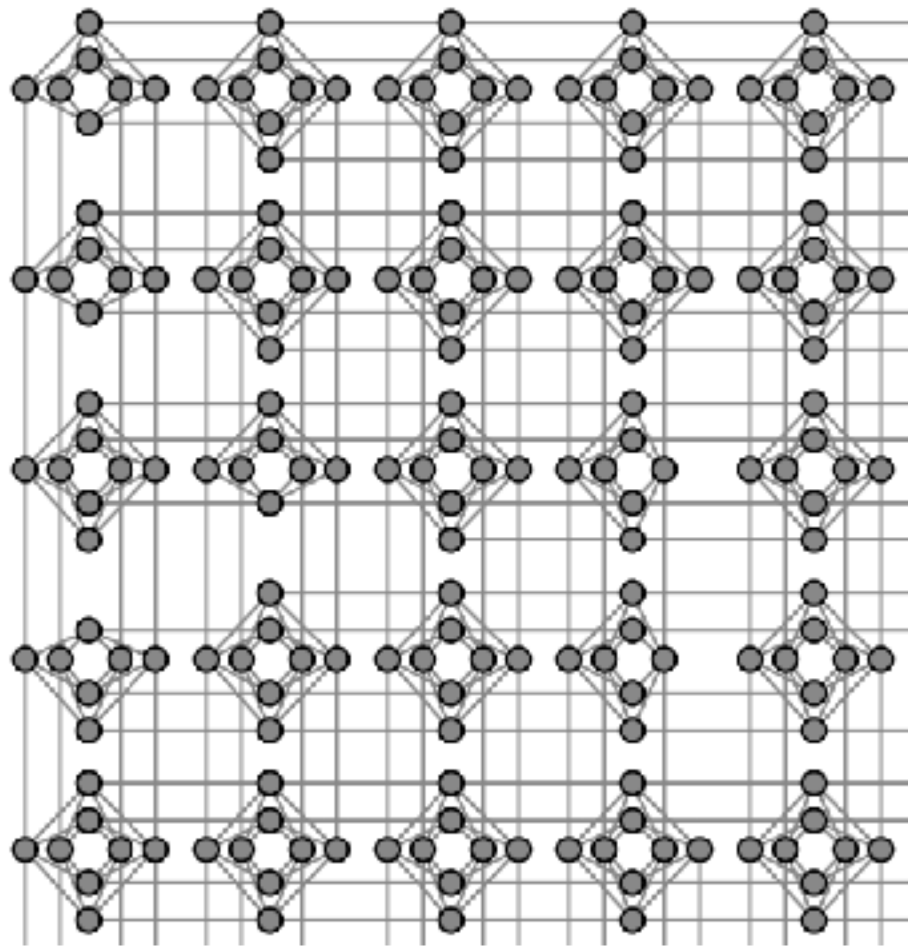

$$\mathcal{C}_5$$

$$2^{194} \approx 10^{58}$$

Detailed Benchmarking Study

Part 1

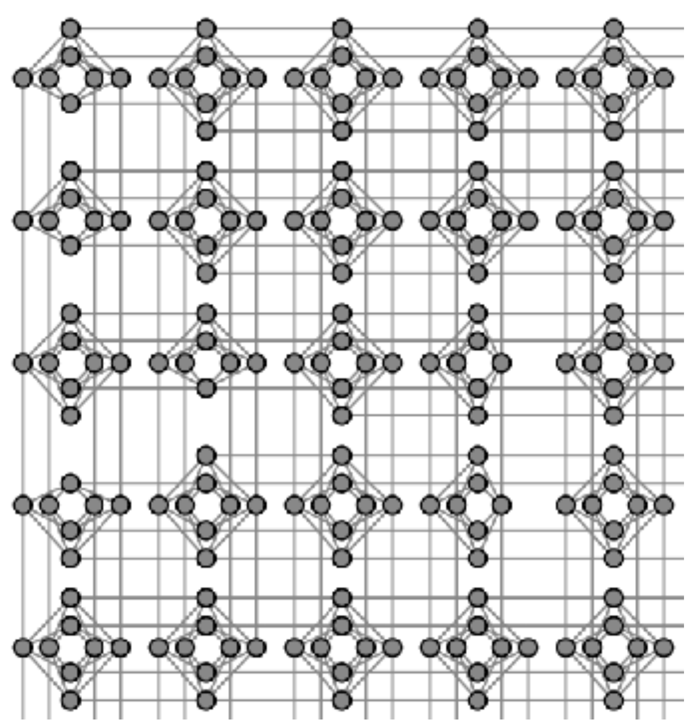

$\mathcal{C}_5$

Quality Validation
(Known Global Optimum)

Wishful Extrapolation

Part 2

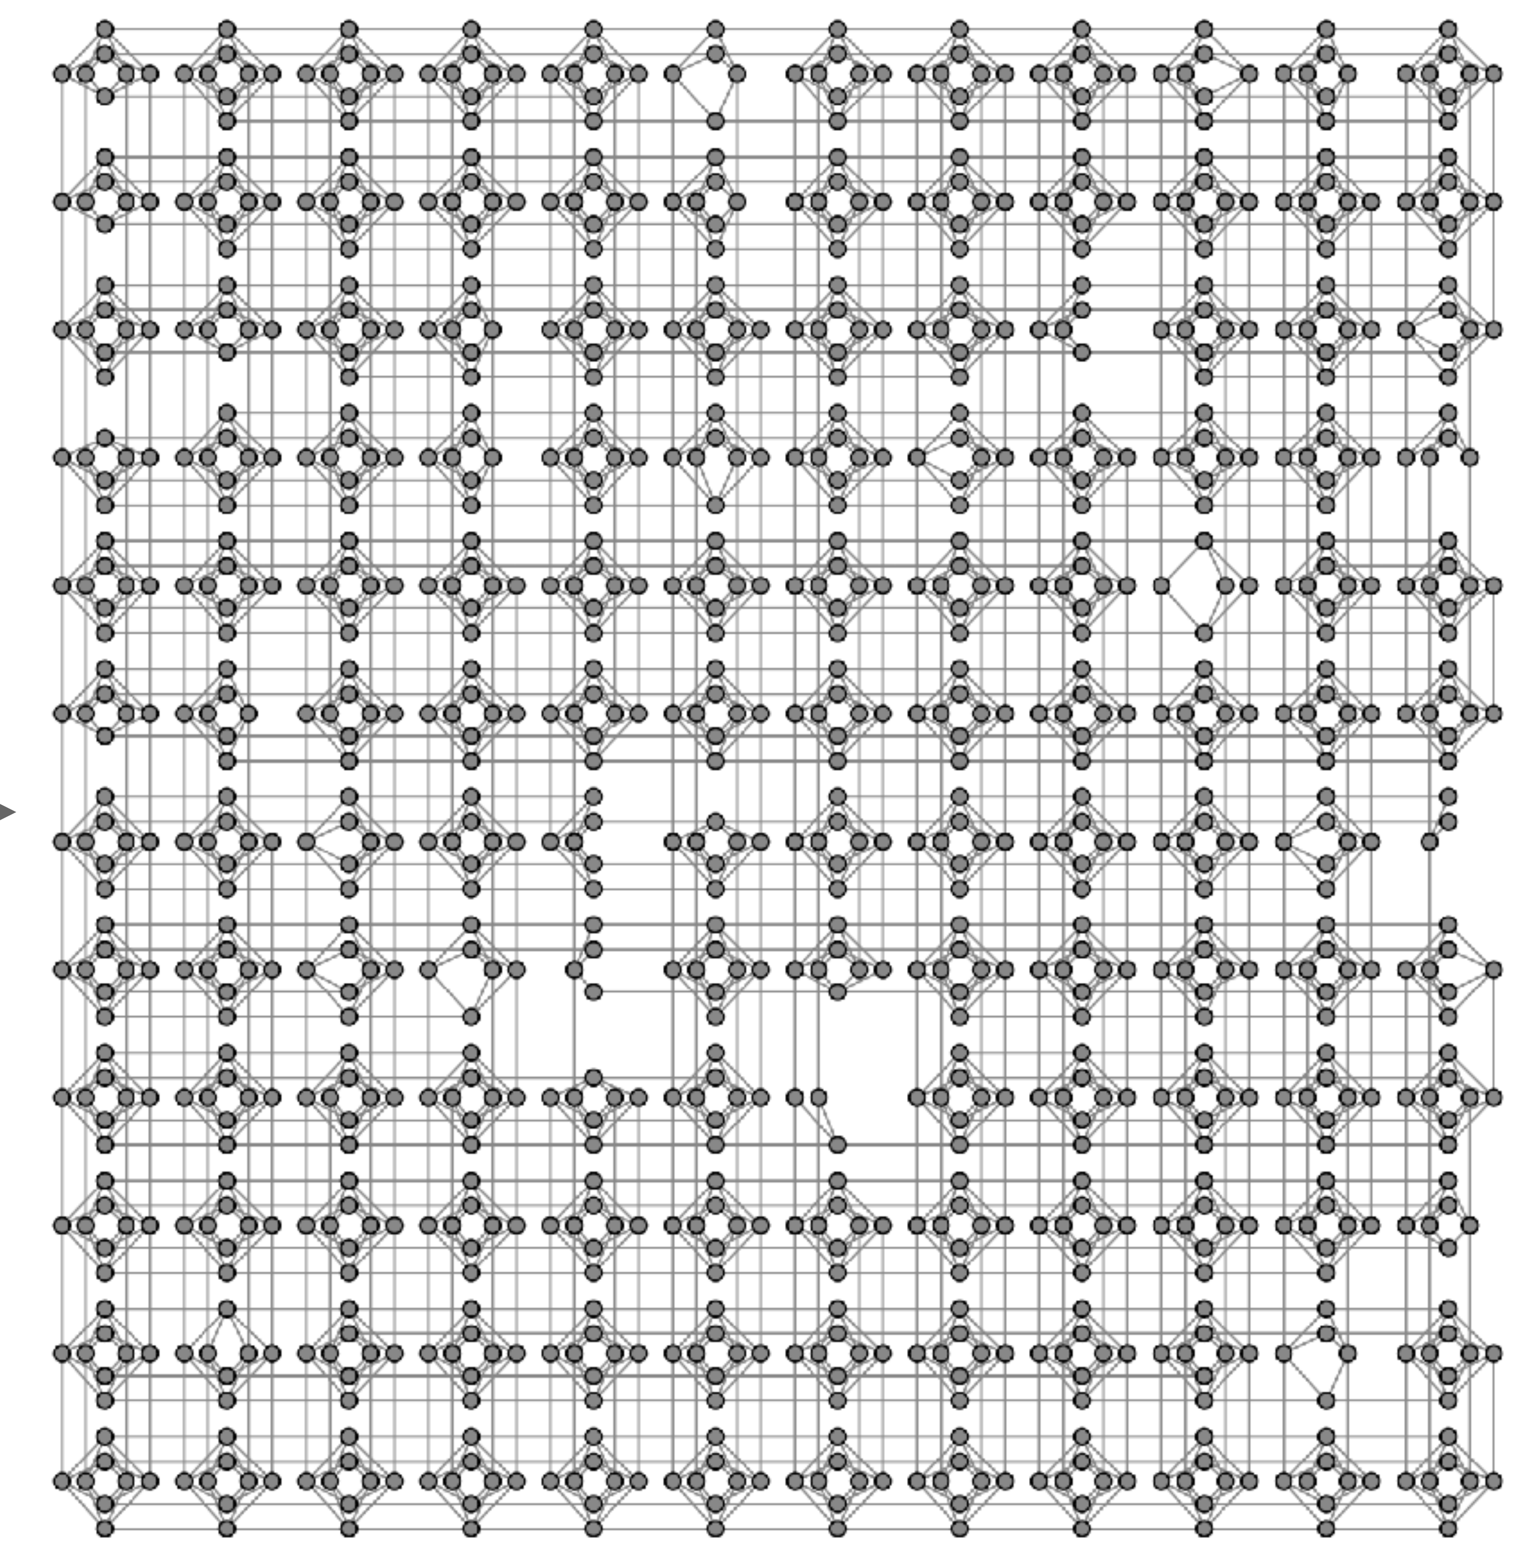

$\mathcal{C}_{12}$

GUROBI OPTIMIZATION

D-Wave

HFS

Detailed Benchmarking Study

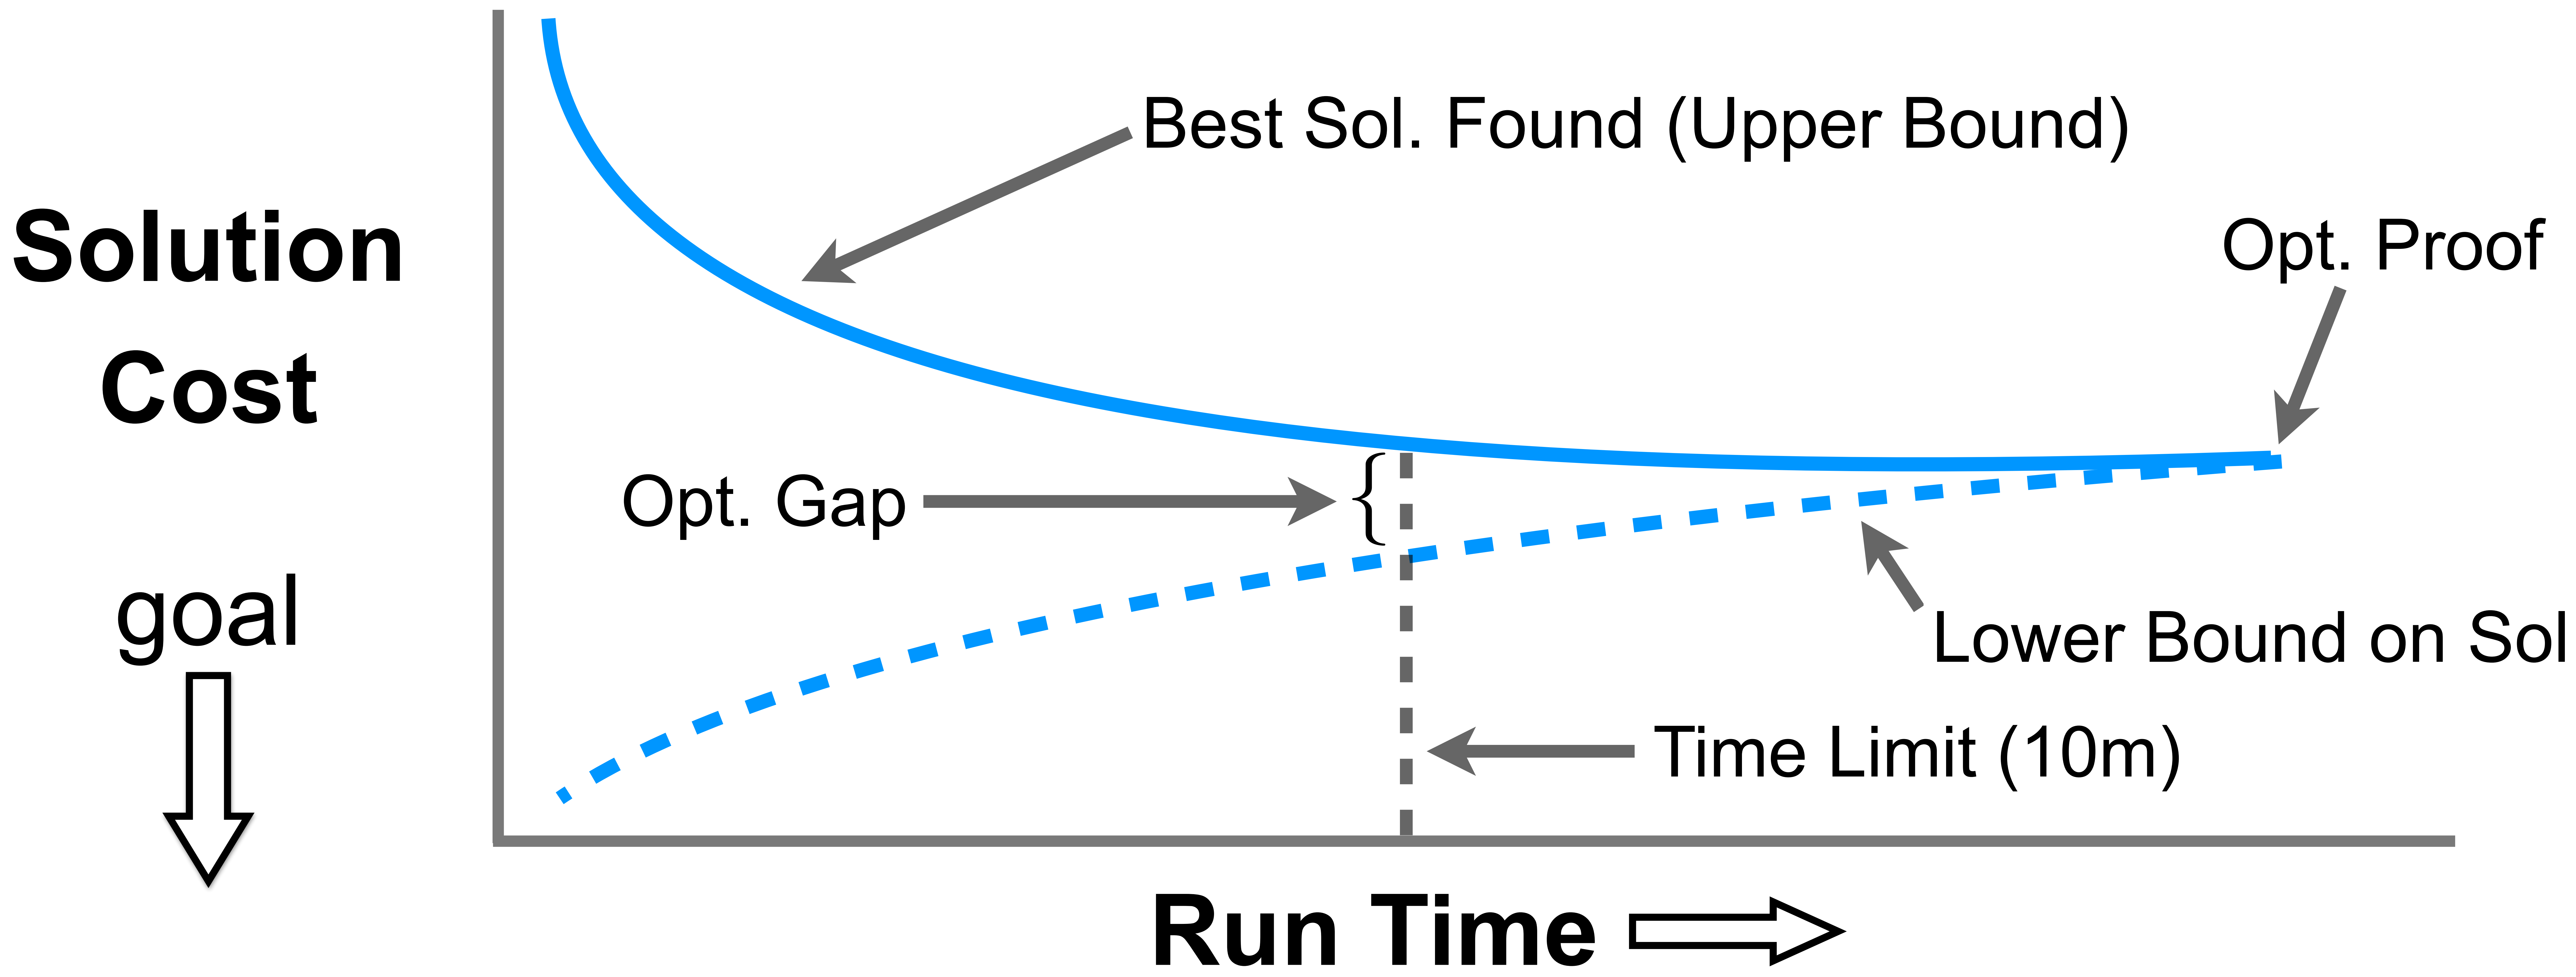

Detailed Benchmarking Study (RAN-1, C_5)

Average Objective / Energy of 200 Cases

Variance in Cases **not** the algorithm

Logarithmic runtime scale

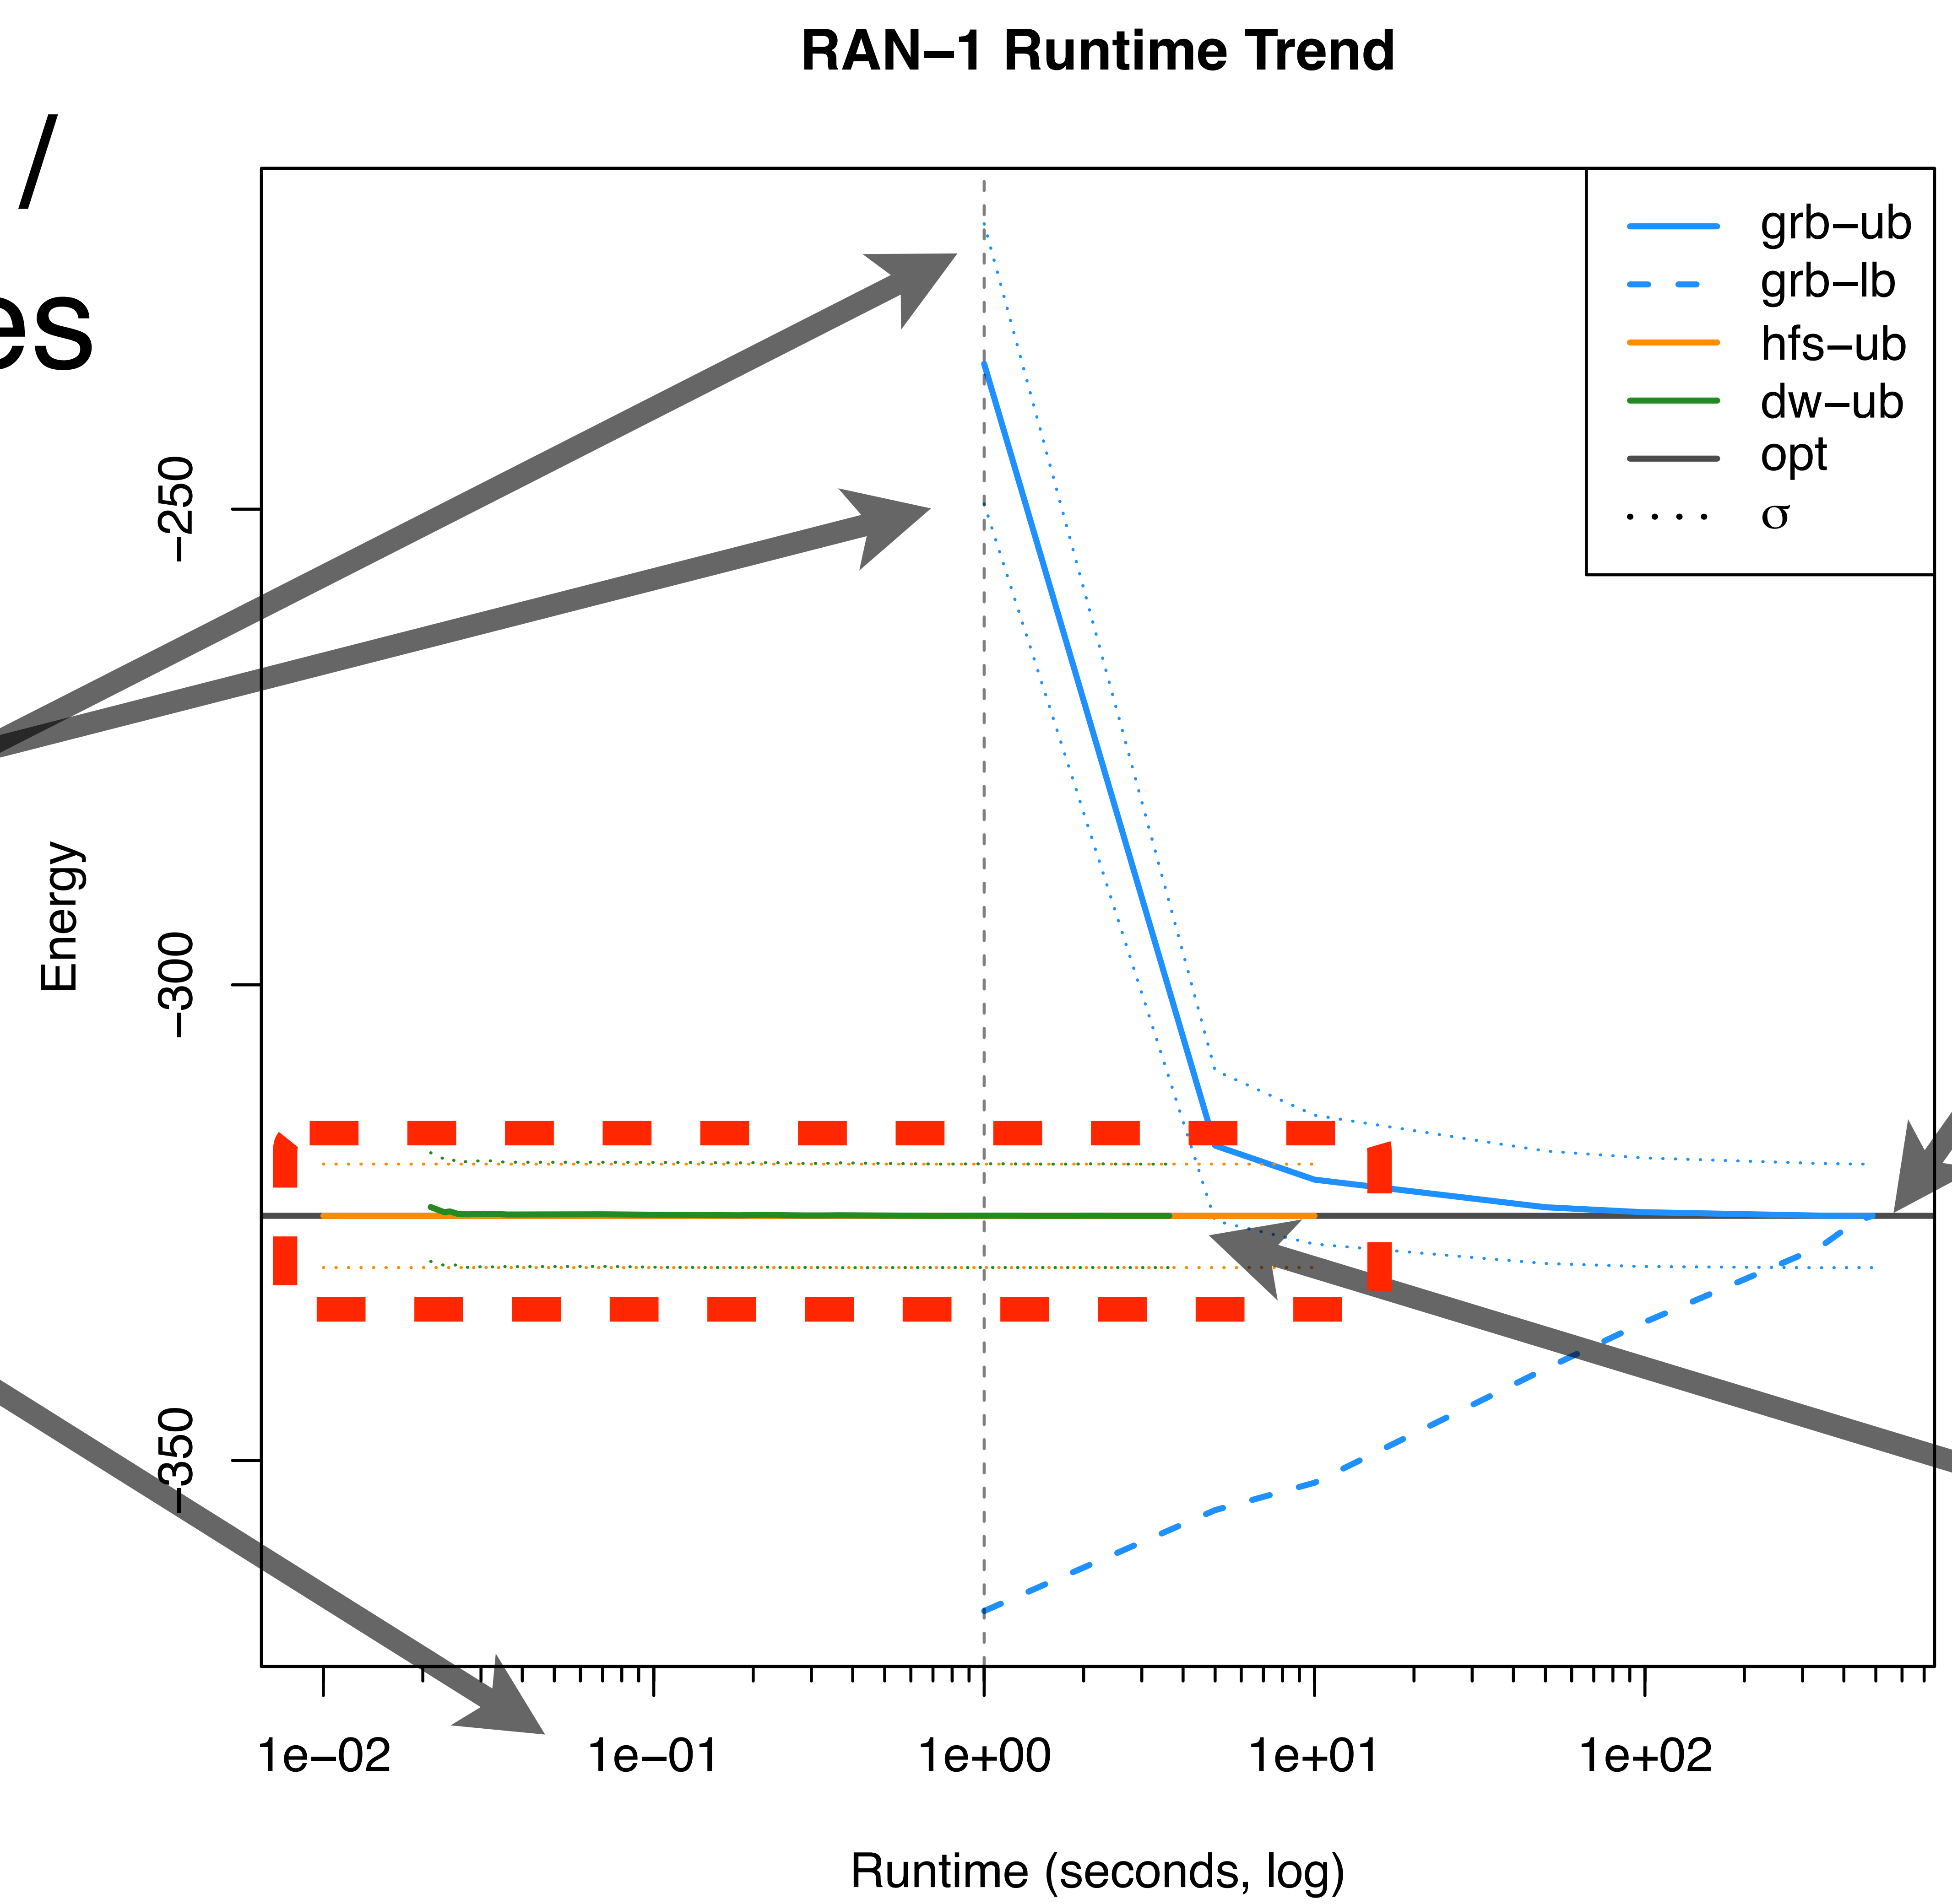

Optimality Proof for all 200 Cases

HFS and D-Wave Indistinguishable

Detailed Benchmarking Study (RAN-1, C_5)

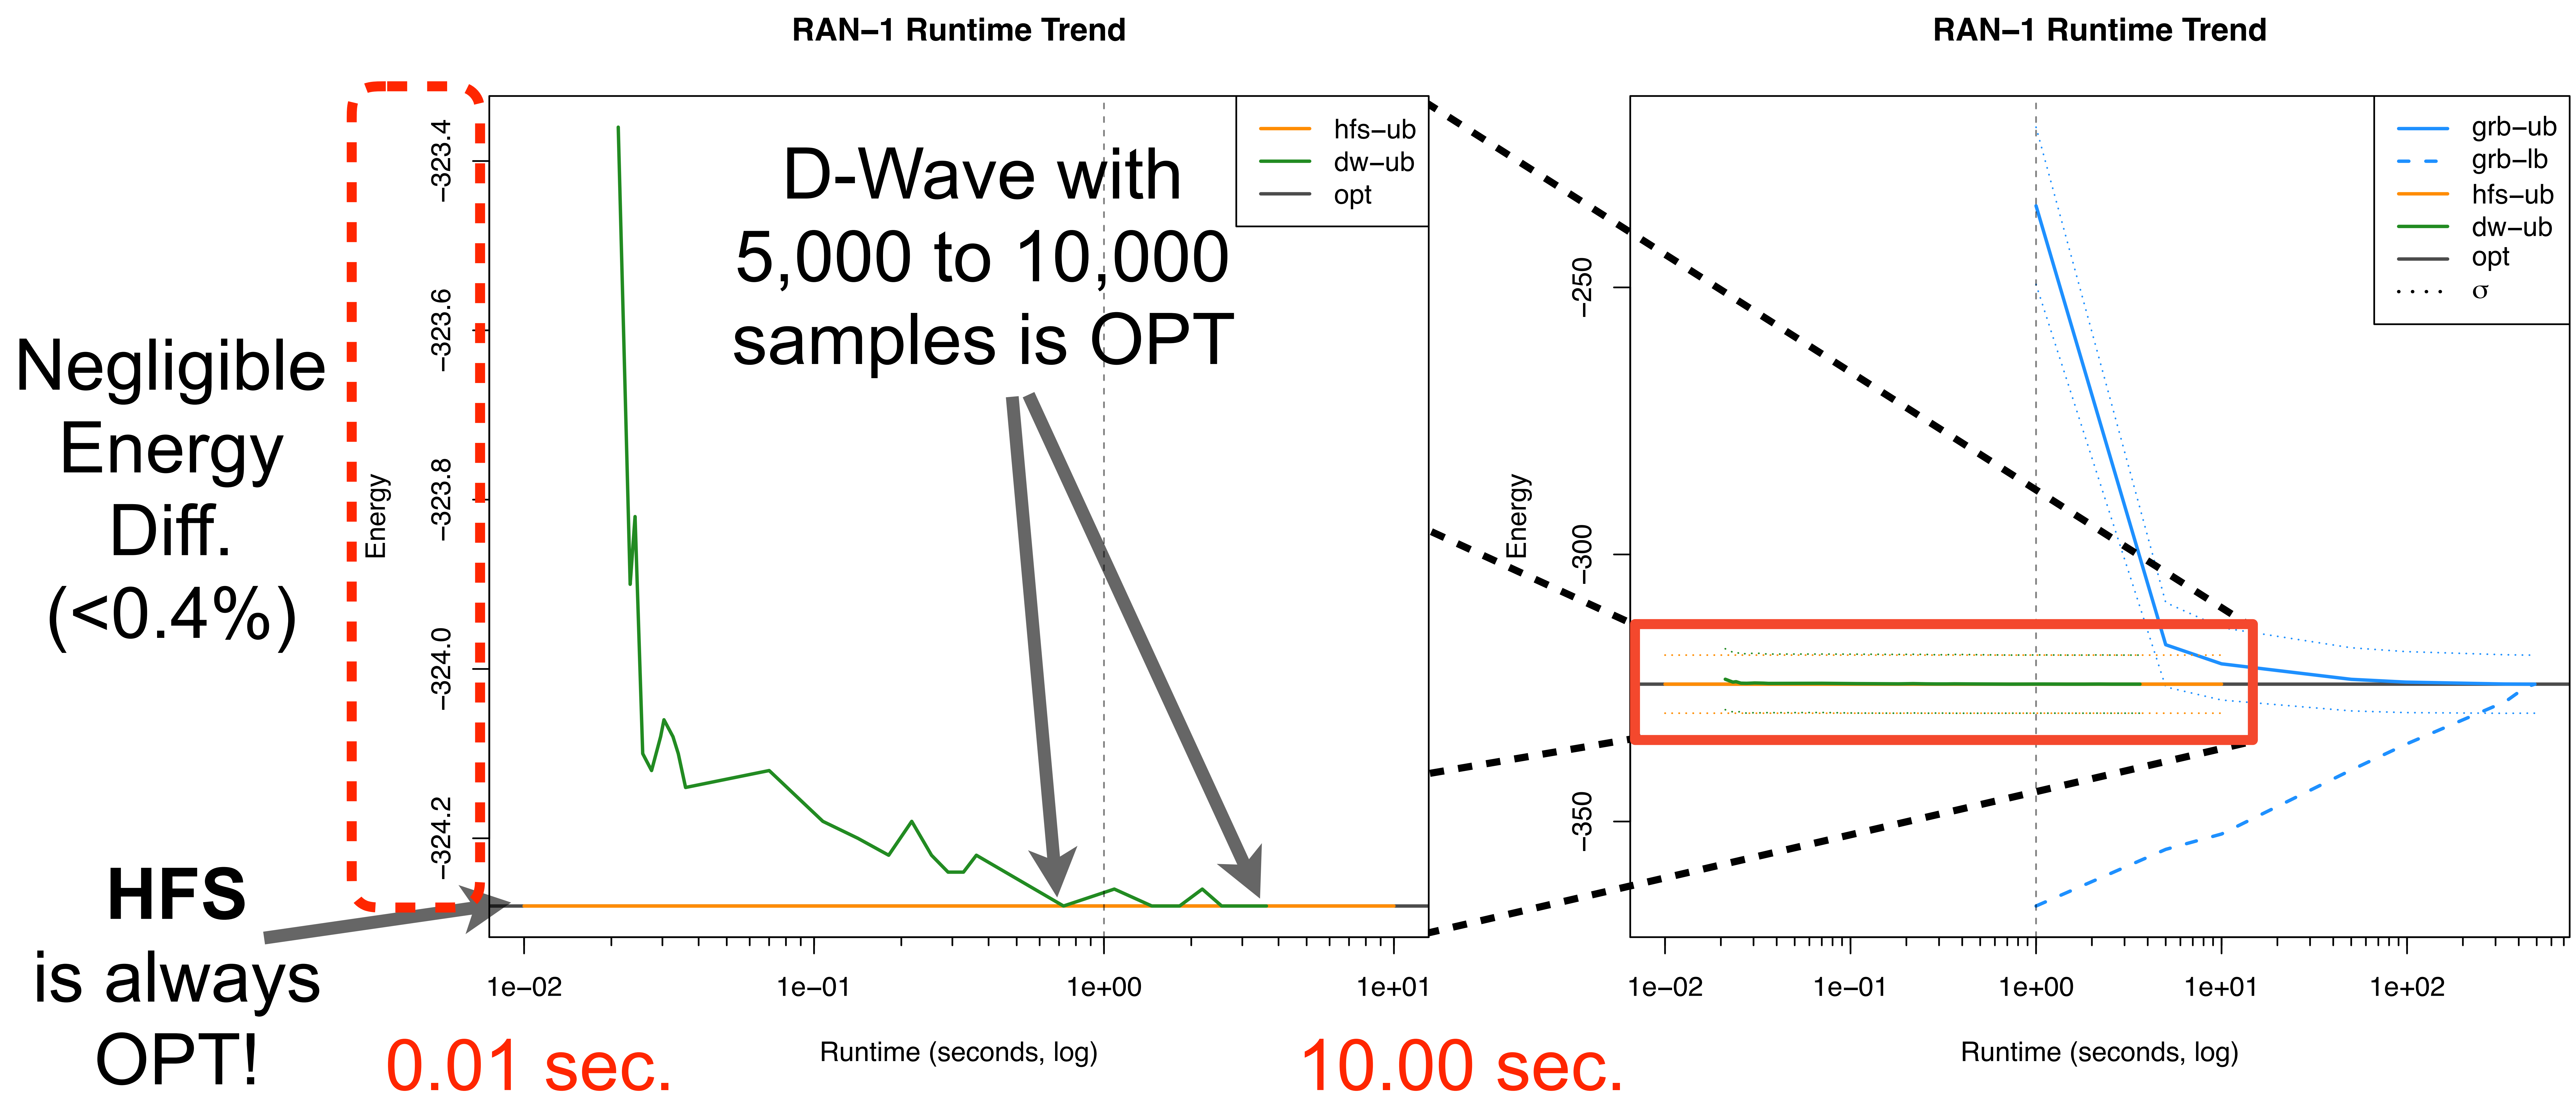

Detailed Benchmarking Study (RAN-1, C_{12})

Negligible Energy Diff. (<0.3%)

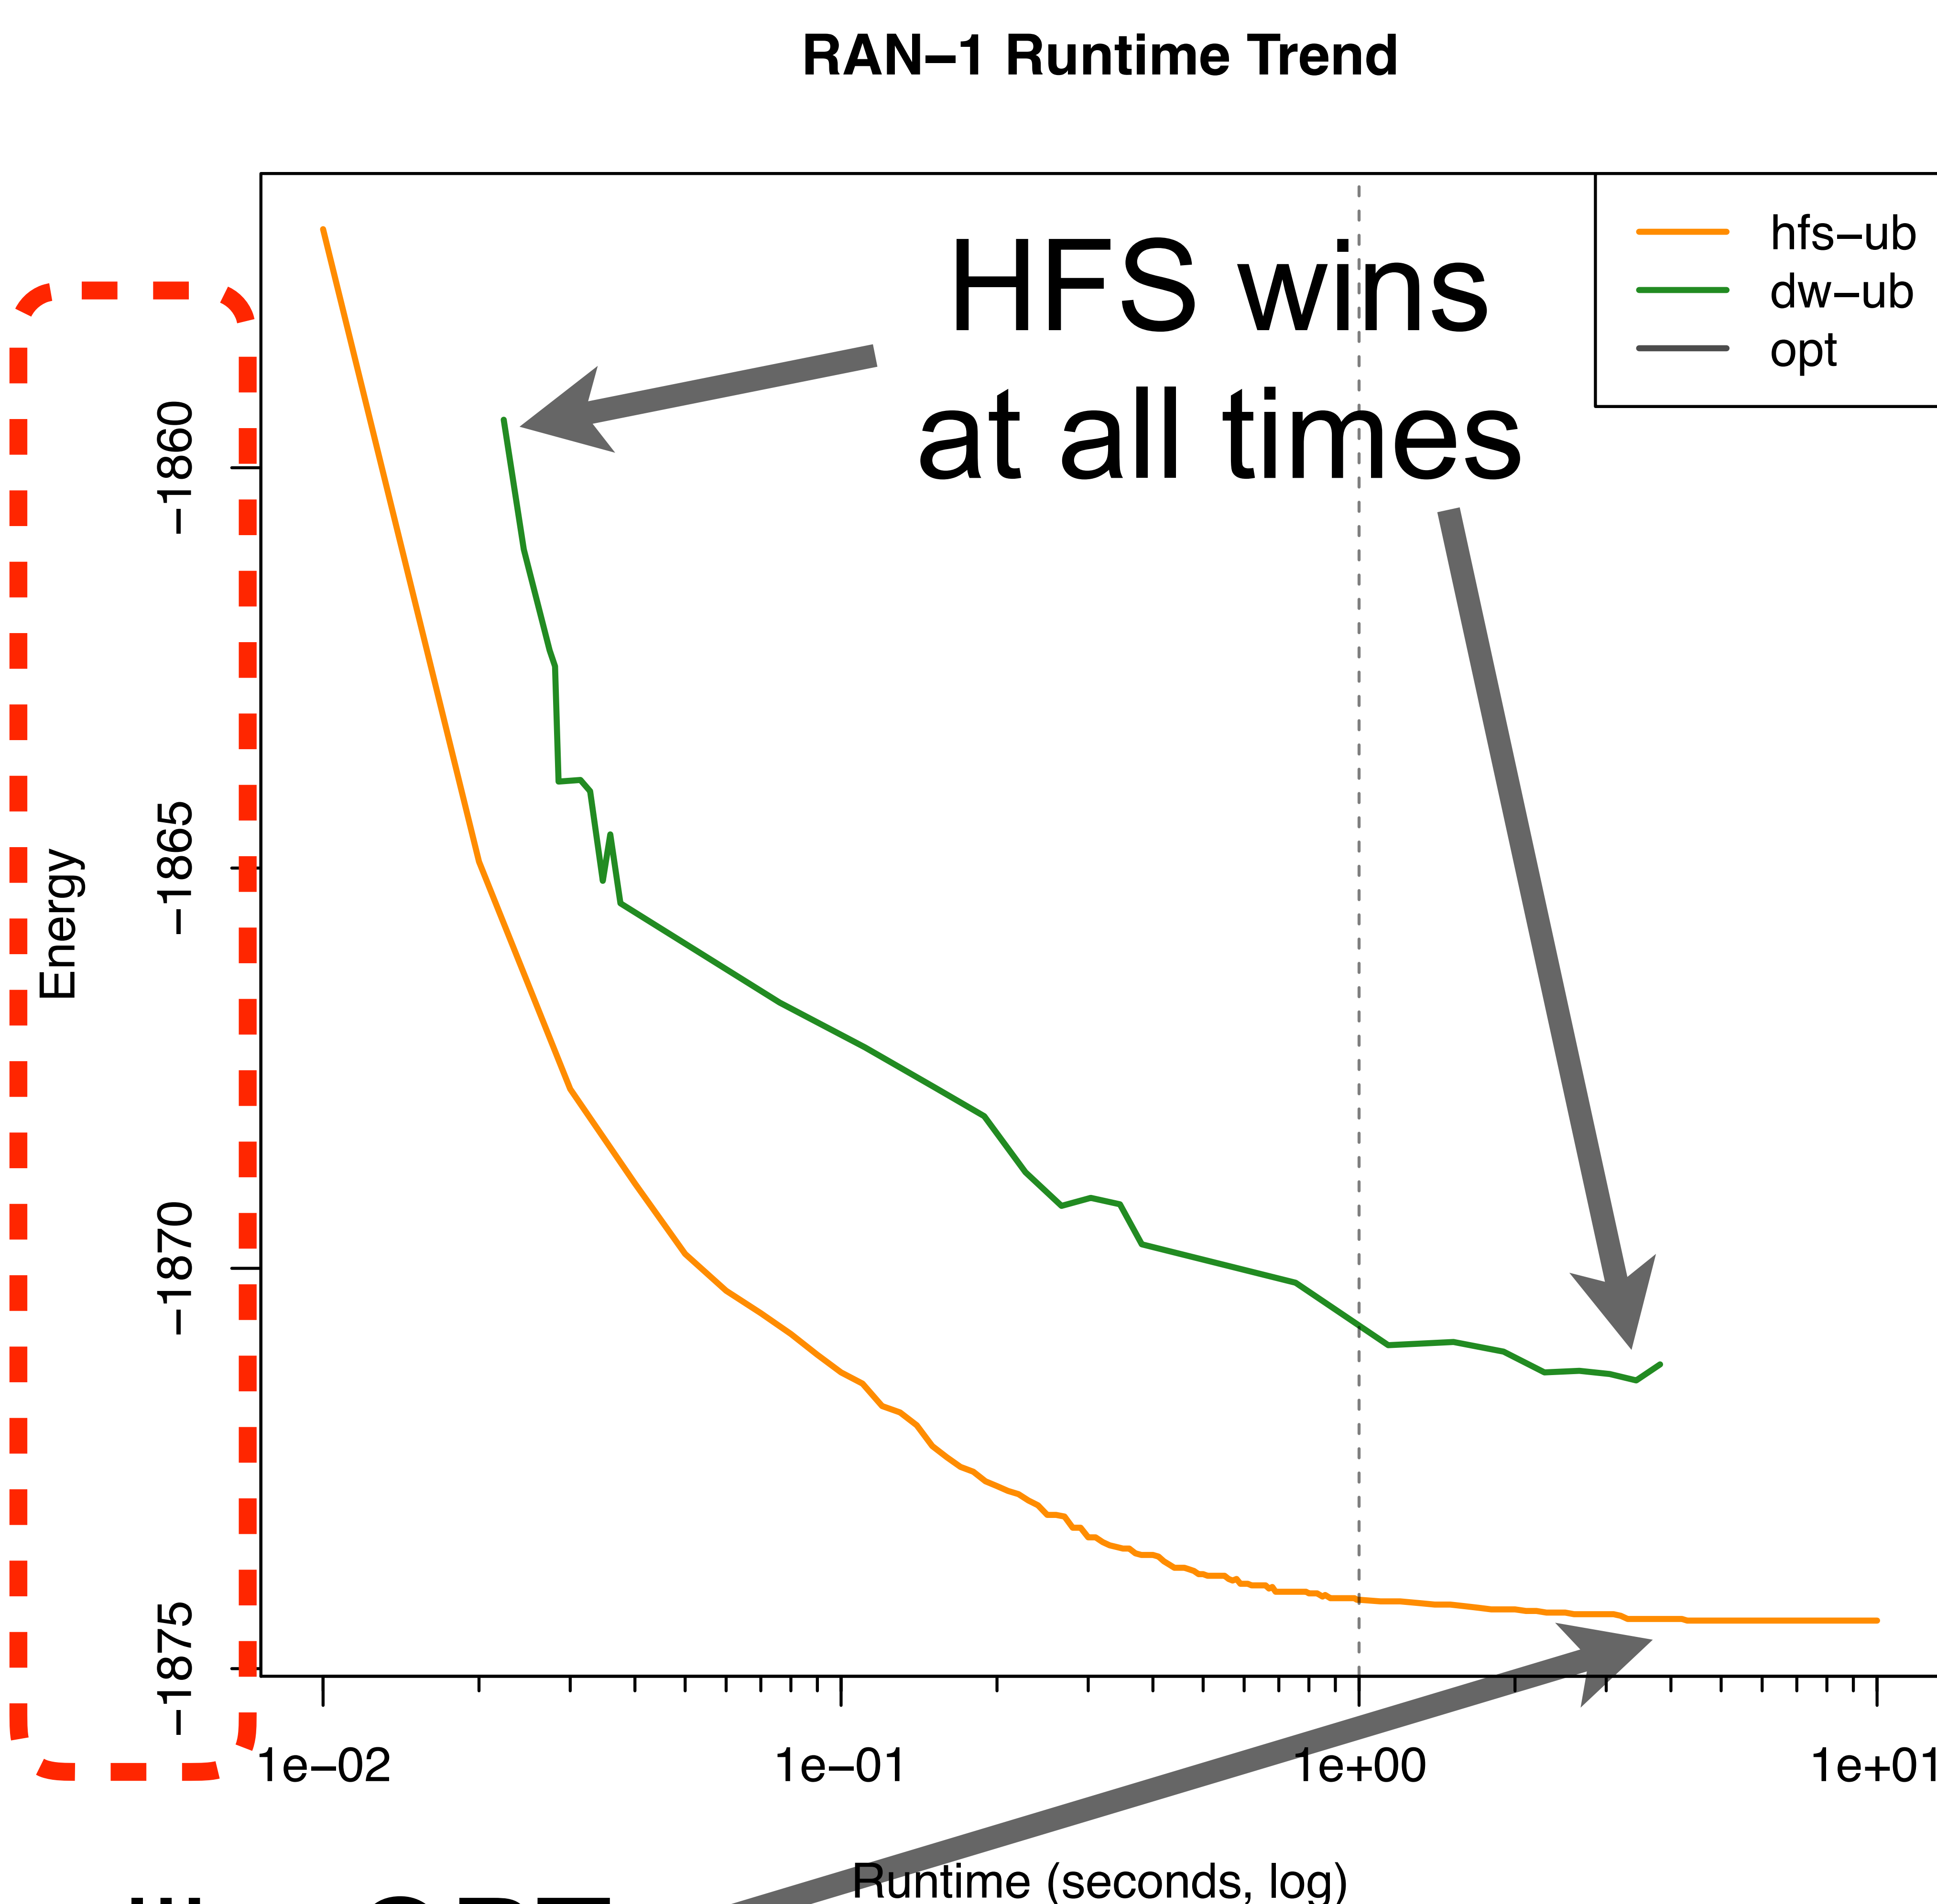

Looks like OPT

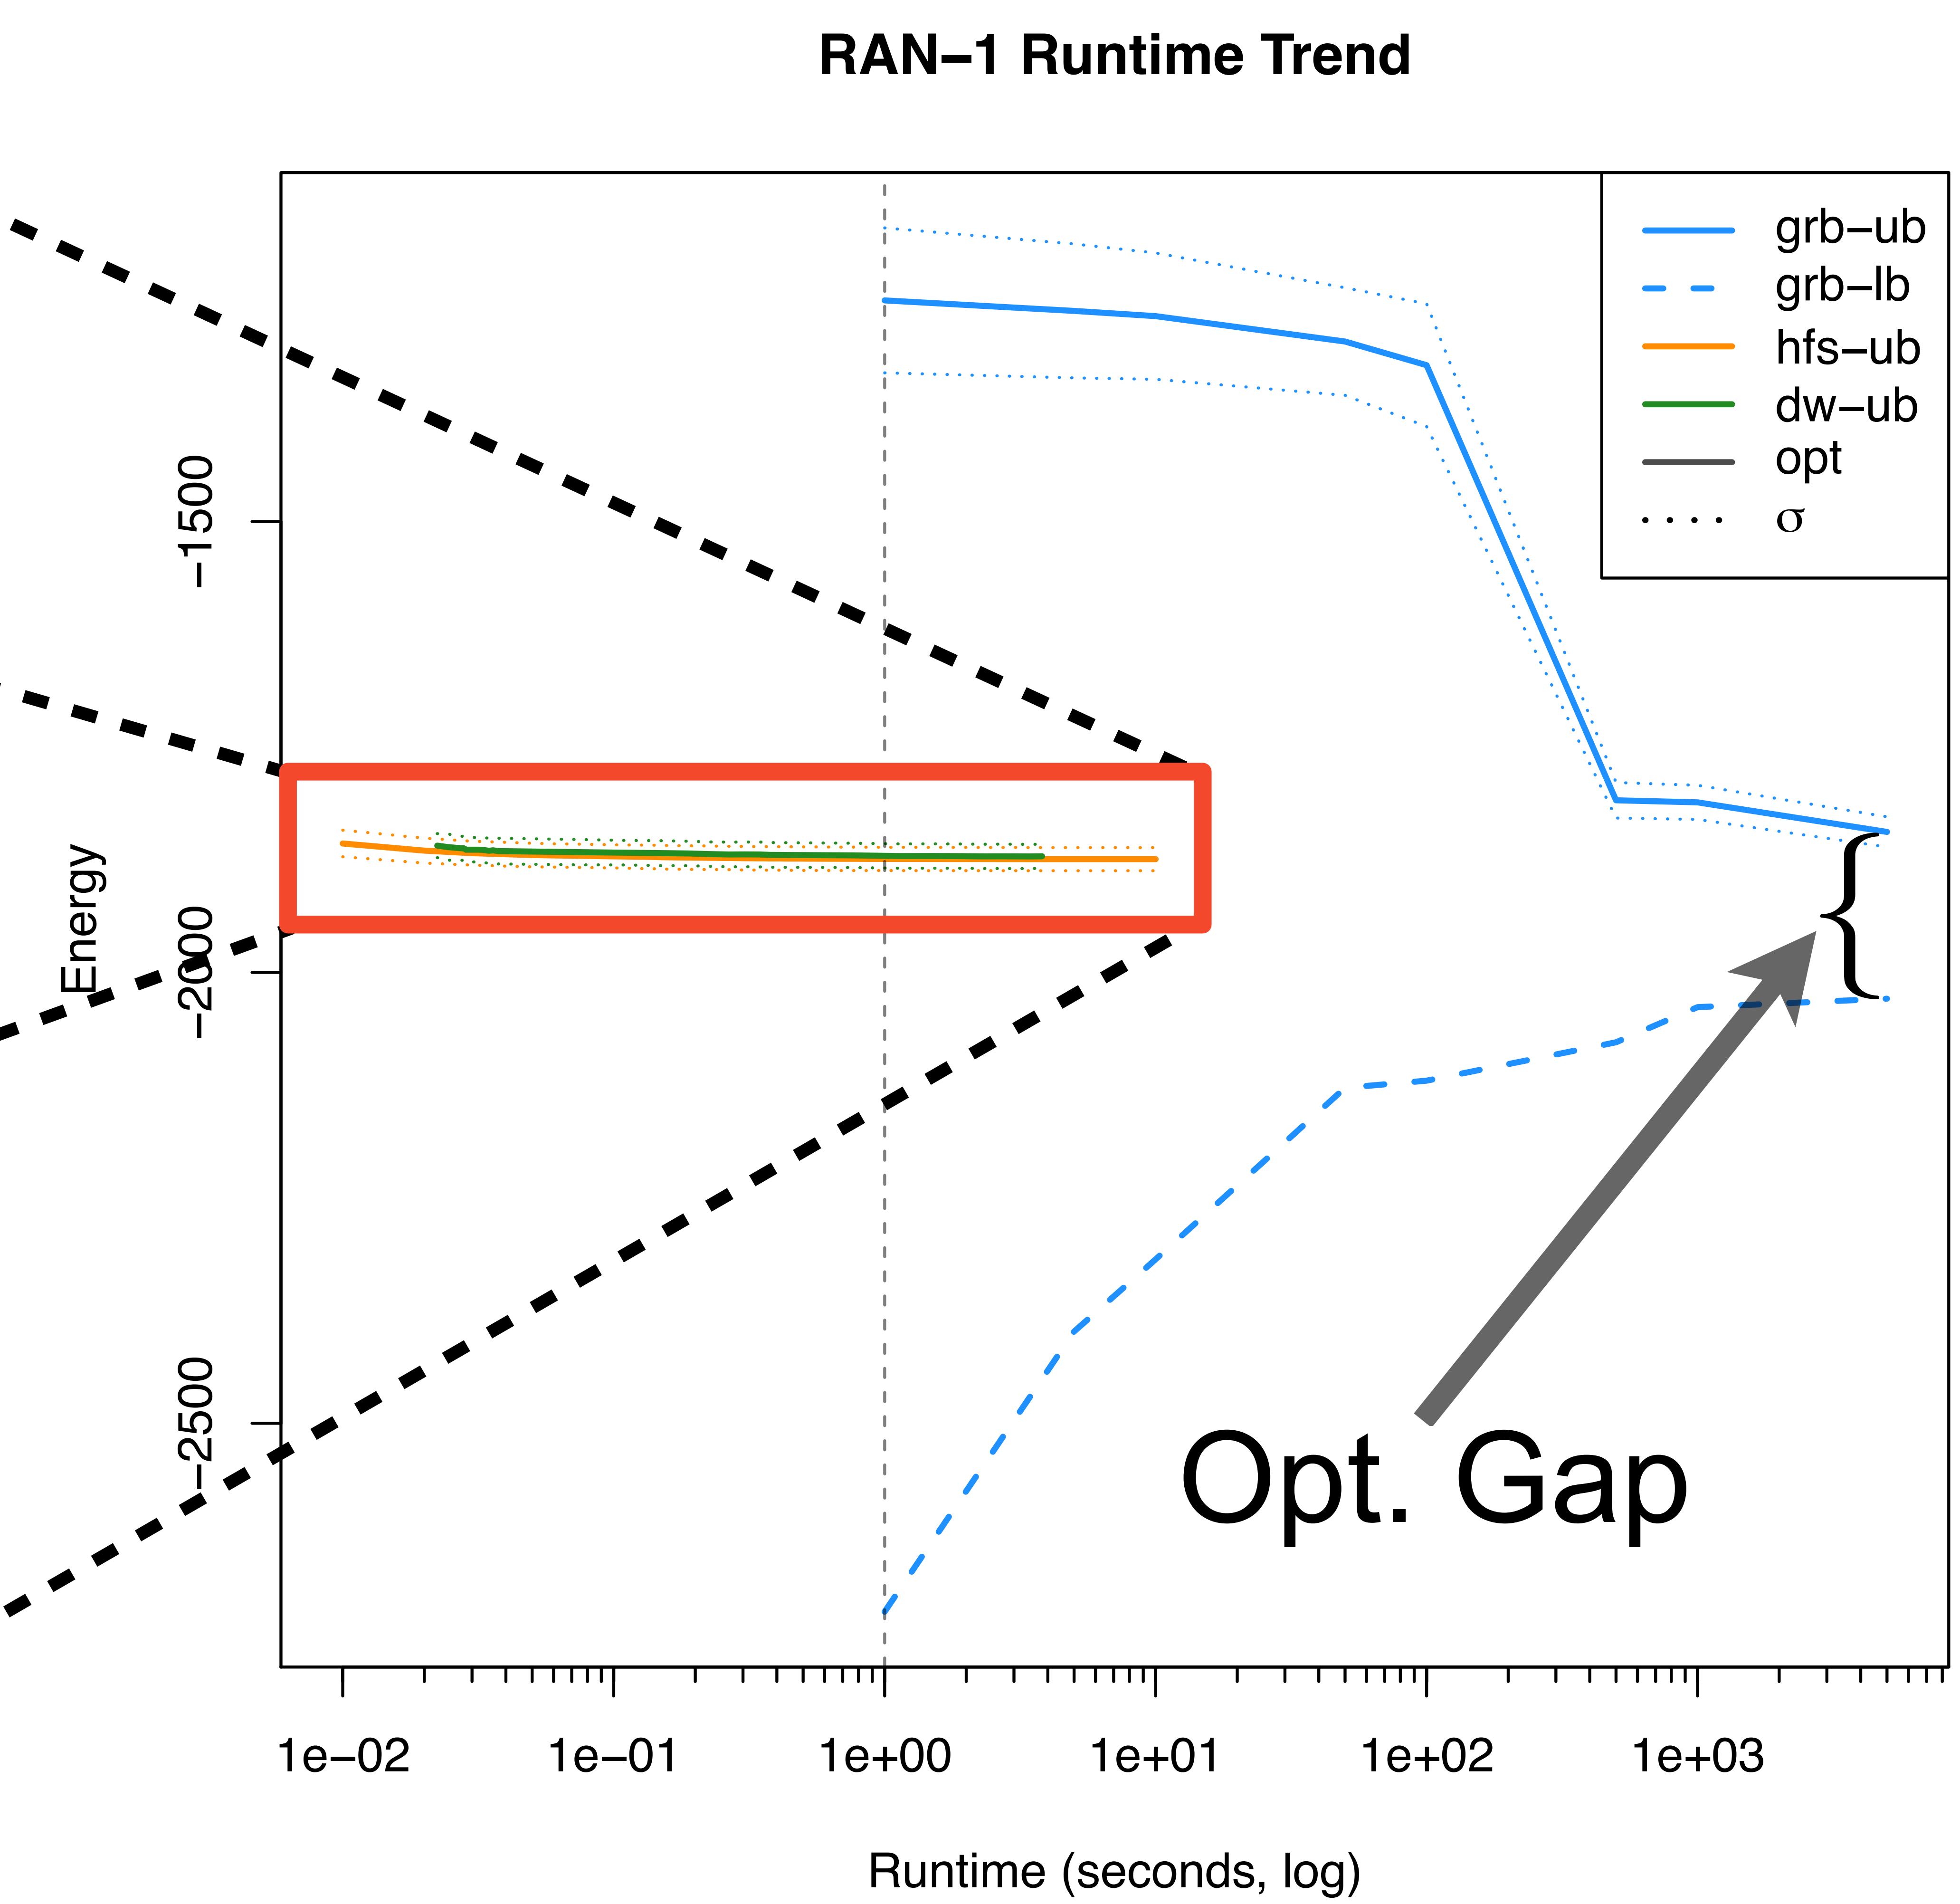

Concluding Thoughts

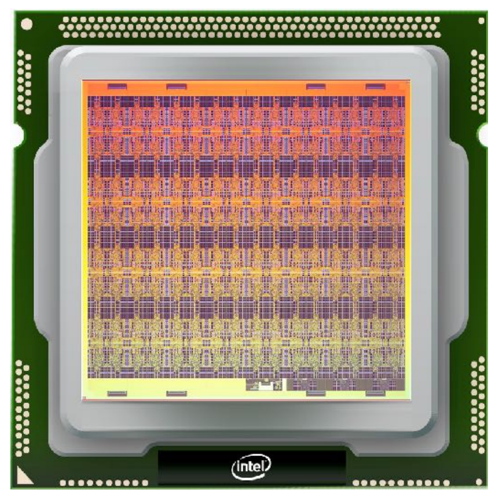

$$\sigma_i = 2x_i - 1$$

QUBO

$$\min : \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij} x_i x_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_i x_i$$

s.t.

$$x_i \in \{0, 1\} \; \forall i \in \mathcal{N}$$

Ising

$$\min : \sum_{i,j \in \mathcal{E}} \boldsymbol{c}_{ij} \sigma_i \sigma_j + \sum_{i \in \mathcal{N}} \boldsymbol{c}_i \sigma_i$$

s.t.

$$\sigma_i \in \{-1, 1\} \; \forall i \in \mathcal{N}$$

$$x_i = \frac{\sigma_i + 1}{2}$$

Ising Processing Units:
Potential and Challenges for Discrete Optimization

https://arxiv.org/pdf/1707.00355.pdf

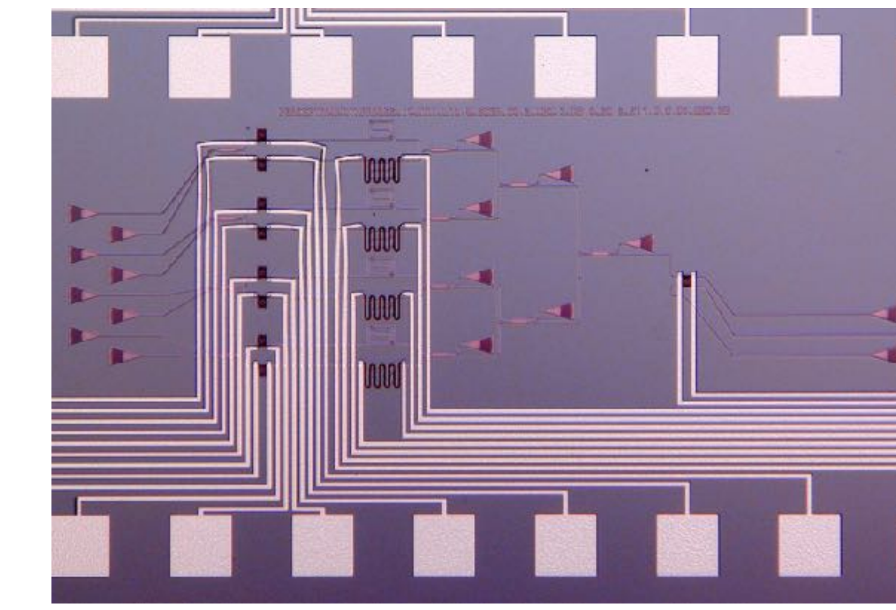

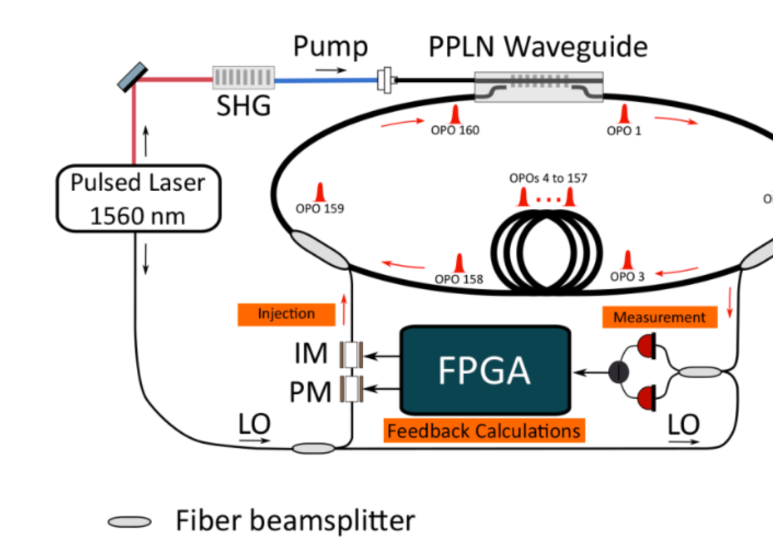

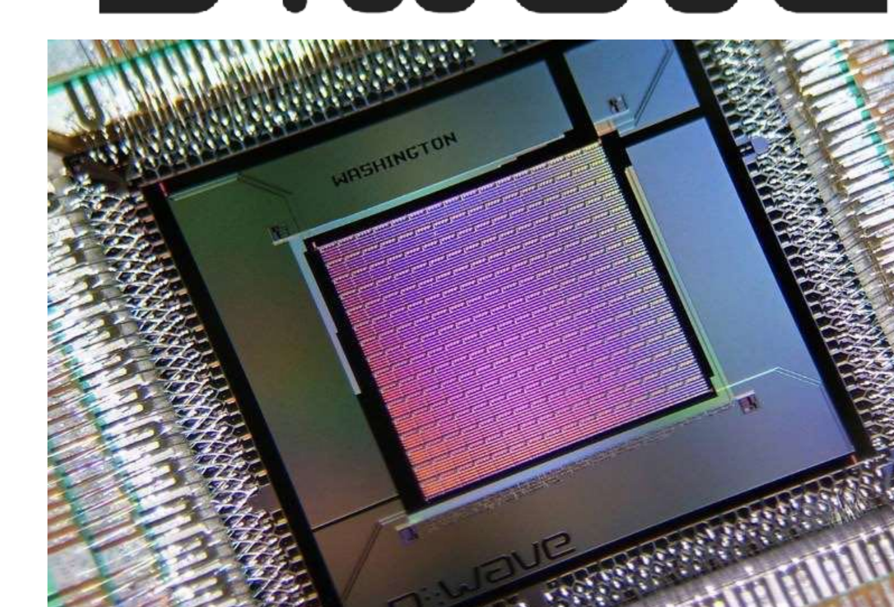

Concluding Thoughts

- **Emerging computational hardware may fundamentally change optimization algorithm development**
  - Ising Coprocessors are like ALUs or GPUs for optimization!

- **Based on initial benchmarking efforts D-Wave hardware is not outperforming state-of-the-art alternatives, but seems competitive**
  - I believe we are on the cusp of a performance breakthrough

- **Watch out for other emerging hardware platforms!**

Thanks!

Questions?